

Try the latest **KDE 4.2 & Slackware 12.2**

Rs 100
ISSN 0974-1054

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 07 ISSUE: 1 March 2009 116 PAGES ISSUE# 74

ForYou

Highly Available

reverse-proxy,
courtesy Heartbeat

D-Bus
Inside-Out
The smart, simple,
powerful IPC



Kernel Virtual Machine

Virtualisation, the Linux way

Flirt with **Perl**
5 rules to follow while coding

Manage Your
Music
Tools that help you
get organised

KDE 4.2

The Answer for Your Desktop

Published by EFY—ISO 9001:2000 Certified

India INR 100
Singapore S\$ 9.5
Malaysia MYR 19



FOSSolution Powered...

- organisations can fight recession
- e-gov projects benefit stakeholders
- students better equipped for job market
- academia ideally placed to go for the win

**Life as an administrator
Is complicated enough...**

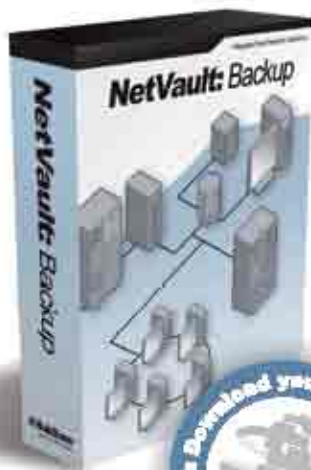
**But Backup and Recovery for your
Linux Servers does not need to be.**

NetVault: Backup simplifies backup and recovery
without compromising **Functionality** and **Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

Permanent FREE Use Edition for Linux is available for download at <http://www.bakbone.com/nvbu/redhat/>



NetVault is a true Linux data protection solution featuring:

- ☐ Online backup
- ☐ Point and click recovery
- ☐ Fully automated protection and recovery support for Linux based applications
- ☐ Virtual Tape Library (VTL) with disk staging
- ☐ SAN support with LAN free backups.
- ☐ Backup duplication for off-site storage of backups
- ☐ Support for MySQL native replication, restore to table level and to alternative databases or instances.

BakBone
The Power of Simplicity™

For more information, please contact
✉ : IndiaSales@bakbone.com
☎ : +91-11-42236155

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 36 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Email Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC501: Programming with OpenGL

RHCE Certification Training

RH033: Red Hat Linux Essentials
RH133: Red Hat Linux System Administration
RH269: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCEB Certification Training

RH8333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RH5429: Red Hat Enterprise SELinux Policy Administration

**New Courses
RH142 & RH184**



**Visit us at - Open Source India
Tech Days 2009, 12 - 14 March
Chennai Trade Center, Chennai**

**RHCE Exams in Bangalore
on 6, 13 & 26 March 2009**

**RH301 from 9 & 23 Mar '09
RH5333 from 7 Mar '09
RH423 from 14 Mar '09
RH5429 from 21 Mar '09**

**Linux Support & Solutions
Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS**

**For more info log on to:
www.linuxlearningcentre.com
Call: 9845057731 / 9343780054
Email: info@linuxlearningcentre.com**

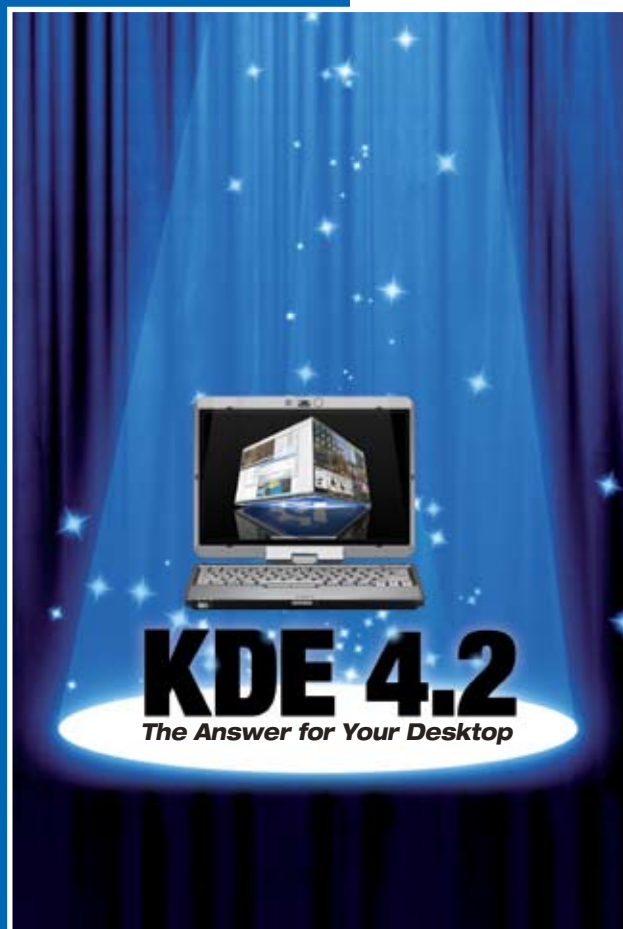
**RED HAT Training Partner
RHCE Examination Centre**



Linux Learning Centre Private Limited

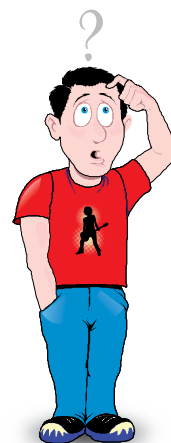
Registered Office & Corporate Training Centre
#635, 6th Main Road, (Adj. Bank of India) Hanumanthnagar, Bangalore 560019
Tel: +91.80.22428538 / 26600639 / 26610999 TelFax: +91.80.28000639
Call: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com
URL: www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore
1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070
Tel: +91.80.26712828



FOR YOU & ME

- 18 Gee, I Like Your Desktop!
- 24 Stop Wasting CDs, Install Linux Straight from an ISO
- 26 Crazy Commands
- 28 Managing Music Efficiently
- 32 Slax 6: Slacks Off To You!
- 35 The GNUified Experience!
- 38 Will FOSS Get Me A Job?
- 40 The Open Movement and the Implications of its Opportunities in Education
- 44 Why Governments Should Adopt Open Source
- 46 Slackware 12.2: Stability Out of the Box
- 50 Open Source: A Panacea for the Recession
- 52 A Matter of Recession
- 94 A Peek Into the WWW, Courtesy MozillaCamp



ADMIN

- 70 KVM: Virtualisation, the Linux Way
- 75 Building A Highly Available Nginx Reverse-Proxy Using Heartbeat

DEVELOPERS

- 53 Watch Out for the Signals!
- 66 Let a Thousand Languages Bloom!
- 84 Flirt with Perl



CONTENTS

GEEKS

- 56 D-Bus: The Smart, Simple, Powerful IPC
- 62 Programming in Python for Friends and Relations, Part 11: Secure Communication
- 80 Building A Server From Scratch —Part 2: Firewalls, Port Forwarding, NAT, DHCP and TFTP
- 86 Lynx: Old, But Still Fresh

COLUMNS

- 83 The Joy of Programming: How to Detect Integer Overflow
- 92 CodeSport
- 98 A Voyage To The Kernel: Day 9, Segment 2.3



REGULAR FEATURES

- 06 Editorial
- 08 You Said It...
- 10 Technology News
- 16 Q&A Section
- 88 Industry News
- 96 Tips & Tricks
- 106 Linux Jobs
- 108 FOSS Yellow Pages



LFY DVD: Slackware 12.2

An operating system with focus on simplicity and stability.

+

Linux is as easy as A B C

A computer-based training programme brought to you by Red Hat.



LFY CD: KDE 4.2

This release adds many new features to KDE4, including some that were notably present in KDE3 but lacked in KDE 4.0 and 4.1, besides some brand new features.

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgroove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhyd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No, 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 24942538, 24925651,
24927383 Fax: 24950392
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear Readers,

The BIG news is the release of Lenny (Debian 5), which is perhaps the best Valentine's Day gift a free software enthusiast could have expected. Debian releases have always been special, particularly because you do not see a new 'stable' release from this project every semester or so, unlike most others. Although many do run the 'unstable' or 'test' versions, every 'stable' version still merits special attention because, as a popular Slashdot quote goes, "They don't release things until you have to fire rockets at the thing to stop it from working."

We, at LFY, are overwhelmed by the amount of enquiries we have received about when we shall be including Lenny with the magazine. Although many of you might have looked forward to it being bundled with the March issue, we could not—unfortunately—because it was released in the middle of the month, and although we were aware of the previous announcement of the release date, we were not sure if it will go live for certain and therefore could not plan it in. So, guess you have to wait for another month. But hey, I think we can make up for the disappointment. This month's DVD has Slackware, which is another of the few distros that likes to concentrate only on things that are 'stable'. You can certainly bet your mission-critical stuff on it.

KDE 4.2 was another significant release, and the awesome desktop experience has left many of us speechless. Err... correction... one of our team members seems to have been overly verbose rather than tongue-tied while filing a report on it! Anyway, we hope you also enjoy the release as much as we did. Try it out from the Live CD included this month.

Moving on, Linux on the mobile (and mobile devices) has been all over the news in February. Well, it has actually been in the news for more than a year now. But, I am talking about the number of announcements and releases made last month, especially with respect to Android, followed by LiMo—a website like LinuxDevices.com pulls out

around 20-30 updates. This includes lots of announcements from Chinese manufacturers; rumours on Dell also entering the mobile phone space with an Android-powered device; the delay in the release of Motorola's Windows Mobile-powered phones, and plans to concentrate on Android; besides HTC's second generation Android phones to be released in Europe this spring (hmm... what about India? :-)).

Hopefully, we will get to hear the buzz from India on this front during the dedicated seminar that has been planned as part of Open Source India (OSI) Tech Days 2009. The seminar, titled 'FOSS-powered mobile phones and devices' is scheduled for March 13, and despite being a half-day session it seems to have the highest number of registrations

Debian 5 is perhaps the best Valentine's Day gift a free software enthusiast could have expected.

amongst all other seminars planned during the conference.

I must thank all of you for the great support in shaping up OSI. It is turning out to be an event that we are all looking forward to. Hope to meet many of you at the Chennai Trade Centre on March 12, 2009.

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com



a GNU/Linux Operating System exclusively available in
localised Indian languages
Hassle-free and quality protection
against virus, piracy and spyware

a hope a creation an innovation
every Indian's mounting dream



BOSS GNU/Linux, an Indian operating System

BOSS (Bharat Operating System Solutions) is a debian based GNU/Linux distribution developed by C-DAC, Chennai in order to benefit the use of Free/Open Source Software throughout India. BOSS GNU/Linux is a key deliverable of NRCFOSS promoted by Department of Information Technology (DIT), Ministry of Communications and Information Technology (MC&IT), Government of India. BOSS GNU/Linux consists of pleasing desktop environment coupled with Indian language support and other packages that are relevant for Indian users and government domain. Subsequent version will support the educational domain as well.

BOSS GNU/Linux Desktop version 3.0 is available in a single DVD with Install, Live and Utility options. Currently it supports 18 Indian Languages: Assamese, Bengali, Bodo, Gujarati, Hindi, Kannada, Kashmiri, Konkani, Maithili, Malayalam, Manipuri, Marathi, Oriya, Punjabi, Sanskrit, Tamil, Telugu, Urdu.

BOSS GNU/Linux advanced Server has unparalleled features that includes user-friendly GUI Front end and supports Intel and AMD x86/x86-64 architecture.



Support Centers

Bangalore: 080-28523300 • Chennai: 044-22542226 • Hyderabad: 040-23150115
Kolkata: 033-23573950 • Mohali: 0172-2237054 • Mumbai: 022-26201488 • New Delhi: 011-24301313
Noida: 0120-3063344 • Pune (H.O.): 020-25694093 • Tiruvananthapuram: 0471-2314412

You said it . . .



There's a great demand for Debian 5, as expected. :-)



I regularly subscribe to LFY and am eagerly awaiting the Debian Lenny DVDs. Could I request that the complete set of DVDs be included in forthcoming issues?

—**Padhu, Pollachi**

ED: *The complete DVD set of Debian comprises of 5 DVDs this time—so bundling the complete set is a little difficult. However, we will definitely bundle DVD 1 with our April issue. We realise that DVD 1 includes most of the essential software for most types of users. Whatever is missing (which won't be much, unless a user has some specific needs) can be downloaded from the Net. The other option is to bundle the rest of the DVD over a period of months—but that would mean having to pass on other major releases, starting with Mandriva in April. Let's see what we can do.*



I am a regular reader of LFY. I find the CD/DVDs of the latest Linux distros bundled with your magazine, very useful. As you have already included Fedora10 in the January issue, I was wondering when you would be including Debian 5. Also, please publish an article with details on setting up a high availability database server using Debian.

—**Chiatanya Kulkarni, Pune**

ED: *It is always heartening when our work gets appreciated by our readers. The topic you've mentioned is very good—we'll certainly try to include the article in the coming month. In the mean time, do let us know how you find the article on Heartbeat, which is included this month.*



The February issue of LFY just rocked. The DVDs were good and I am thrilled with LFY for

providing the latest Linux OSs. The articles on RMS, Metalinks and on Distromania were very educative. I would like to get even more information about Metalinks.

Also, can we get a complete article with detailed steps on how to compile a kernel from scratch? Not like the *Linux Kernel in a Nutshell* book written by Hartman. I need an explanation that helps me download the latest kernel and then guides me on compiling it. I have tried it many times but have not succeeded.

Finally, I am happy to report that I have seen some improvements in the quality of the LFY DVDs and CDs. Before signing off, could I request that the March issue carries all the 5 DVDs of Debian 5.0 (Lenny)... The Debian Etch DVDs had some issues while copying them. Please see to it that this problem is addressed.

—**Ananth Gouri, by e-mail**

ED: *It's great to hear that you liked the content of our February issue. The Metalink Wikipedia page at en.wikipedia.org/wiki/Metalink lists an overwhelming list of applications that support it. As for an article on compiling a custom kernel, although we've carried many articles on the subject in the past, I guess it's about time we took a look at it from the point of view of general users. We should be able to arrange something in the forthcoming issues.*

And well, when you get the Debian DVD in April, you can rest assured that the quality will be top-notch.



First, the bad news: when I tried to visit www.openitis.com mentioned in LFY, I got a pop-up from my AV. Now the good news. I have recently moved to LFY from another technical magazine and am very pleased to see the quantity of content, at such an economical price. The content is very easy for different user

groups—my wife and sister read the magazine whenever time permits. LFY is a magazine that even my friends wait for me to pass on to them.

Having worked in the IT sector for more than 11 years, I still have that urge to keep testing the bleeding edge operating systems and applications. Keeping only the media I have received from the Ubuntu online site, I have distributed LFY and its DVDs to almost 50 people (including students) by now. I have even shipped media to colleagues in my company, across the country.

I am sure you'd be bundling Debian Lenny sooner or later, however, please can you ship BackTrack and/or System Rescue CD images in one of your forthcoming issues?

A great magazine by a great team! Keep them coming!

—**Mitesh Vohra, by e-mail**

ED: *We feel proud to know that LFY has managed to satisfy both yours and your friends' open source needs. Speaking of the AntiVirus popup, thanks for bringing it to our notice. We've already passed on the information to the Web team, and I believe it has been taken care of.*



I have recently started reading your magazine. The articles on OpenVAS and GlassFish were informative. Please publish some articles on PostgreSQL and MySQL database connectivity, in addition to something on how to configure Sendmail on RHEL 5. This will help us learn how to configure Linux as a mail server. I was able to set up a security server on Linux in the recent past. So, next up, we're thinking of setting up a Linux-based mail server.

—**Anand Nayyar, Ludhiana**

ED: *Well, thank you for your feedback. We're glad that you find the content of the magazine useful in*

You said it...



your work and hope you continue to deploy more and more open source software in your organisation. We'll definitely try to include articles on DB connectivity and Sendmail in the forthcoming issues of LFY.



I've been a subscriber since LFY's very first issue (Feb 2003). First, I would like to express my appreciation for all the efforts of LFY/EFY team, the readers, authors and Linux lovers.

I wanted to enquire about two things:

- Fedora Core 10 DVD seems to work only on new (the latest configuration) PCs. I have run the DVD on a DELL (Pentium 4) system with 1 GB RAM, Intel 82845G/GL graphics card, and also on an IBM (Pentium 4) PC. Fedora seems to have an issue with the display on both the systems—for example, dialogue boxes are not displayed properly. But, if I run the DVD on a latest PC, I could see all GUI dialogue boxes running successfully.
- I did not find any 'anniversary news' in the February 2009 issue. Am I missing anything?

—**Naresh Bhalala, Patni**

Computer Systems Limited

ED: Thanks for your feedback. Fedora 10 works fine on our P4 system with an 845 MoBo. In fact, we have even tested it on pretty old Celeron systems and didn't encounter the issues you're facing at our end. Since it's an official Fedora 10 ISO, you should definitely post your query at the official Fedora forums. This will also ensure the developers take note of the issue, besides, there's a better chance of you getting a workaround to the problem. As for the 'anniversary news', it's the "Leader of the Free World" feature—the exhaustive interview, was really the highlight of our anniversary. :-)



I started reading LFY from Dec 2008. I am a B. Tech student (6th semester) and am currently using Fedora. I liked the article "What's in the Glass(Fish)?—Part 2: Getting Started with the Application Server", by Rajeev Kumar. The problem is, I unfortunately couldn't get the Jan 2009 issue of LFY because the stock finished in the city. Can you please mail me the PDF of Part 1.

—**Bharat Chand, by e-mail**

ED: Looks like everyone has words of praise for the GlassFish article. As for the PDF version of the requested article, it must have reached your inbox by now, we hope :-)



As a regular reader of your magazine, I wanted to make a suggestion. In your earlier issues you used to feature case studies of companies or organisations that use Linux in their IT infrastructure, such as Breach Candy Hospital. Could you please start this series again so that it will be helpful to everyone, while learning how organisations are implementing Linux.

—**Brahmaji Rao C, by e-mail**

ED: Point noted! We'll definitely try to include case studies of Linux deployments whenever we can in the upcoming issues.



It feels good to write back to you after a long time, even though I am a regular reader! First of all I would like to wish the LFY team a belated new year as well as hearty congratulations for publishing another fantastic issue. Just wanted to know how you see the future of Linux, and OSS in particular, given the frequent cost-cutting measures taken by organisations (both in terms of technology and manpower)?

Also there's a little request from my side. Why not start a forum

in the magazine as well as in the website, where people could come up with queries or ideas to develop an application and others who are interested in developing/contributing to it, can come forward. I suggest this in the best FOSS spirit and hope it will help people stay tuned even when the times are rough.

—**Sreekanth Narayan, by e-mail**

ED: Although, we're all hit by the economic meltdown, we believe the worldwide recession could turn out to benefit OSS. In fact, open source is one of the things that could help us get out of these hard times. As for your suggestion about a forum, we completely agree with you there. Since an online version of the magazine is on the cards, placing a forum there is also under consideration. Let's hope we can launch the site in the next few months.



Greetings and congratulations for an astounding issue :) The anniversary issue had everything that I had been pestering you guys for since quite some time :-)

The presentation and content is just top notch... I loved everything. The new interface of regular sections except for 'Know How' was just too good. The designers deserve some applause. :-)

—**Shashwat Pant, Chandigarh**

ED: Thanks, it's great when our readers notice the finer points that we keep working on to improve the magazine. Your congratulatory words have been conveyed to our designers.

Please send your comments or suggestions to:

The Editor

LINUX FOR YOU Magazine
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020

Phone: 011-26810601/02/03

Fax: 011-26817563

Email: lfyedit@efyindia.com

Website: www.openITis.com



Technology News



Debian 5.0 Lenny, finally released

After 22 months of vigorous development and testing, the Debian Project released the Debian GNU/Linux version 5.0 (codenamed 'Lenny') on February 14. This OS runs on computers ranging from palmtops and hand-held systems to supercomputers, and on nearly everything in between. It officially supports 12 processor architectures—Sun SPARC, HP Alpha, Motorola/IBM PowerPC, Intel IA-32, IA-64, HP PA-RISC, MIPS, ARM, IBM S/390, and AMD64 and Intel EM64T. Debian GNU/Linux 5.0 Lenny adds support for Marvell's Orion platform, which is used in many storage devices. Those supported include the QNAP Turbo Station series, HP Media Vault mv2120, and Buffalo Kurobox Pro. Additionally, Lenny now supports several Netbooks, in particular, the Eee PC by Asus. It also contains the build tools for Emdebian, which allow Debian source packages to be cross-built and shrunk to suit embedded ARM systems.

Debian GNU/Linux 5.0 Lenny includes the new ARM EABI port, 'armel'. This new port provides a more efficient use of both current and future ARM processors. As a result, the old ARM port has now been deprecated.

This release includes numerous updated software packages, such as the K Desktop Environment 3.5.10 (KDE), an updated version of the GNOME desktop environment 2.22.2, the Xfce 4.4.2 desktop environment, LXDE 0.3.2.1, the GNUstep desktop 7.3, X.Org 7.3, OpenOffice.org 2.4.1, GIMP 2.4.7 and Iceweasel 3.0.6 (an unbranded version of Mozilla Firefox).

It also includes Icedove 2.0.0.19 (an unbranded version of Mozilla Thunderbird), PostgreSQL 8.3.6, MySQL 5.0.51a, GNU



Compiler Collection 4.3.2, Linux kernel version 2.6.26, Apache 2.2.9, Samba 3.2.5, Python 2.5.2 and 2.4.6, Perl 5.10.0, PHP 5.2.6, Asterisk 1.4.21.2, Emacs 22, Inkscape 0.46, Nagios 3.0.6, Xen Hypervisor 3.2.1 (dom0 as well as domU support), OpenJDK 6b11, and more than 23,000 other ready-to-use software packages (built from over 12,000 source packages).

With the integration of X.Org 7.3, the X server auto-configures itself with most hardware. Newly introduced packages allow the full support of NTFS filesystems and the use of most multimedia keys out-of-the-box. Support for Adobe Flash format files is available via the swfdec or Gnash plug-ins. Overall improvements for notebooks have been introduced, such as out-of-the-box support for CPU frequency scaling. For more information on the latest release, and in order to download the OS, visit www.debian.org.

Kernel 2.6.28 with Arch 2009.02

The Arch Linux team has released new installation images, dubbed version 2009.02, that comes with: kernel 2.6.28, ext4 support; rescue and maintenance capabilities for ext4 root partitions; fallback ISOs with the ISOLINUX Kernel 2.6.28 Grub-based images; inclusion of AIF (Arch Linux Installation Framework), etc. You can download the images from www.archlinux.org/download. Lead developer Aaron Griffin also announced that their "...goal is to bring out coordinated releases following the rhythm of kernel releases, in order to provide optimal hardware support."

Grab the SimplyMepis 8

Warren Woodford has released SimplyMEPIS 8.0, the community edition of MEPIS 8.0. The new version utilises a Debian Lenny stable foundation enhanced with long term kernel support.

In addition to Linux kernel 2.6.27.18, MEPIS 8.0 includes KDE 3.5.10, OpenOffice 3.0, and Firefox 3.0.6. It also comes with Bind 9.6, while IPv6 is enabled out-of-the-box. Virtualisation can be easily achieved by downloading KVM 84 and libvirt 0.6.0 from the MEPIS 8.0 package pool. Mirrors from where you can download the ISO images of this release are listed at www.mepis.org/mirrors.

GUI installer in Vector 6

VectorLinux 6.0, codenamed 'Voyager', is now available for download. A non-GUI based installer has been a big disadvantage for the otherwise user-friendly desktop. Finally, the sixth version brings forth its very own stable GUI installer.

Additionally, the software repository now hosts over a thousand packages. The main desktop is based on XFCE 4.4.3 with a custom theme and artwork unique to VectorLinux. LXDE is installed as a secondary desktop option. You can visit vectorlinux.com/downloads to download Voyager.



National Conference on Open Source Software

May 25th - 26th, 2009, Mumbai, India

Organised by

C-DAC, Mumbai

Supported by

IEEE Computer Society, Mumbai and Chennai chapters
Computer Society of India, Div II on Software & SIG-OSS

Call For Papers

NCOSS—2009 is a forum to bring together the various groups working on developing Open Source Applications catering to specific domains in the ICT world—education, health, accessibility, localisation, e-commerce, disaster management, expert systems, machine learning, etc. A number of high-quality software solutions are available in many of these areas, for example, SugarCRM, Koha, Drupal, Moodle, Sahana, CollabCAD, etc. Work on these systems require a combination of domain knowledge and development expertise. Much of the public awareness in open source is focussed on desktop, operating system and general productivity tools. With this background, NCOSS-09 has chosen to focus on the layer above this, bringing together groups working on various application domains.

The conference will present experiences in deploying FOSS applications, comparative studies among competing software solutions, efforts in adapting and localising FOSS applications, development of new applications, etc. The conference will consist of the following:

- Invited talks by experts from India and abroad
- Presentation of contributed papers selected based on refereeing by a panel of referees
- Exhibition by industry and academia
- Pre conference Tutorials (on May 24th)
- Panel discussion

TOPICS

Papers are invited on the topics listed below: (Other application areas may also be considered).

Accessibility	Machine Learning and Data Mining	e-Governance
Indian Language Computing	e-Health	e-Commerce
Localisation	Knowledge Management	Disaster Management
e-Learning	Collaboration Technologies	Content Management
Information Extraction and Retrieval		

INSTRUCTIONS

- Papers must report original work carried out by the authors. The work can include enhancing existing Open Source applications for specific requirements, development of new solutions and comparative analysis of competing solutions. Direct survey or overview papers are not acceptable.
- Length should not exceed 10 pages of A4 size in length (approx. 5000 words) including figures, etc.
- Papers should be in English.
- An abstract of about 100-200 words and the area(s) under which the paper can be categorized, must be included with the paper.
- The author names and affiliations along with the main area of the paper should be given only on a separate cover sheet. Papers should be in one of the following formats: PDF, RTF or ODT. Accepted papers will be published in the conference proceedings.

Visit <http://ncoss.cdacmumbai.in> for more details and paper submissions.

Media Partner:



Technology News



OMAP MDP supports LiMo, Android

Texas Instruments has unveiled an enhanced version of its OMAP 3 processor-based development platform—the Zoom™ OMAP34x-II Mobile Development Platform (MDP), which has been designed, developed, and manufactured by Logic. The new platform is targeted at smartphone and mobile internet device (MID) application developers who want to create applications for the Android Mobile Platform, Linux, LiMo, Symbian and Windows Mobile. It offers them a robust handheld form factor with the wireless connectivity technologies, enhanced imaging, video, display technology, software, as well as an optional 3G modem and optional DLP Pico™ projection technology module that enables big picture experiences in the palm of your hand.



Out-of-the-box features of the Zoom OMAP34x-II MDP include the following: 4.1-inch WVGA multi-touch display with a QWERTY keypad in a landscape, handheld form factor; high performance OMAP3430 applications processor that supports up to 720p HD video encode/decode; wireless connectivity technology from TI, including WiLink™ 6.0 (WL1271); a single chip with Wi-Fi, Bluetooth and FM functionality; NaviLink GPS functionality; 8 MP camera sensor; optional 3G modem solution, as well as flexibility to support any third party modem through an extension card. Go to www.ti.com/orderzoom for information on how to order.

ACCESS Linux Platform 3 goes 3D

ACCESS, a provider of advanced software technologies to the mobile and 'beyond-the-PC' markets, showcased ACCESS Linux Platform v3.0 at this year's GSMA Mobile World Congress in Barcelona, Spain. The next generation of the company's flagship mobile Linux platform features advanced UI capabilities and LiMo compliance, and will be made available soon according to the company.



According to the release, "ACCESS Linux Platform v3.0 features an advanced

UI engine and middleware that enable licensees to create state-of-the-art user experiences with Hollywood-style graphics and transition effects with added support for 2.5 and 3D graphics environments. Enhanced flexibility allows different applications, from different environments, to co-exist and be concurrently executed. Content, such as contacts, appointments, videos or photos, can be rendered anywhere, not just within one application."

The ACCESS Linux Platform v3.0 SDK enables the development of native applications for LiMo-compliant devices. By wielding the full power of the Eclipse IDE software development platform, the ACCESS Linux Platform v3.0 SDK runs in the ACCESS Linux Platform Simulator. Also available on the ACCESS Developer Network (ADN) website at www.accessdevnet.com, the SDK allows users to create user interfaces with GUI builder tools.

The new XenServer release is free for all

Citrix Systems has unveiled the new version of XenServer that will be offered free of charge to any user for unlimited production deployment. With this new release, XenServer adds powerful new features like centralised multi-node management, multi-server resource sharing and full live motion. Powerful centralised management enables full multi-node management for an unlimited number of servers and virtual machines; includes easy physical-to-virtual and virtual-to-virtual conversion tools, centralised configuration management and a resilient distributed management architecture. As for Live Motion and Multi-Server Resource Sharing, they incorporate the XenMotion technology that allows virtual machines to be moved from server to server, without service interruption ensuring zero downtime. Also included are optimal initial virtual machine placement and intelligent maintenance mode.

"Free hypervisors with limited functionality have been around for a long time. We see this move as substantially different because it offers a competitive, enterprise-ready virtual infrastructure platform with fully centralised management, live motion and support for unlimited virtual machines and servers—with no strings attached," said Mark Bowker, analyst, Enterprise Strategy Group.

The free XenServer release will be available for download from the Citrix website and other download portals by the end of March 2009. You can preview the XenServer release at www.citrix.com/freexenserver

A stable v 2.6.26-based RT Linux released

The Open Source Automation Development Lab (OSADL) has announced that the 'latest stable' version (2.6.26-8-rt16) of real-time mainline Linux (a.k.a PREEMPT_RT) is now based on kernel version 2.6.26, after successfully testing it in a wide variety of kernel configurations and on many different platforms.

Apart from maintenance fixes, the 'latest stable' version incorporates two significant features: device tree support and improved kernel cache management of the video buffer. According to the release information, the device tree is, "...a (simple) flat data structure containing information about the devices of a given computer board. The device tree source (DTS) is compiled using the device tree compiler (DTC), and the resulting device tree binary (DTB) is integrated into the boot image. The device tree facilitates board configuration and is required for the merging of the two PPC architecture implementations, ppc and powerpc."

The improved kernel cache management of the video buffer, on the other hand, "...makes it possible for the first time to use hardware-accelerated graphics in a real-time system without any side effects of graphics operations on the real-time capabilities of the system. There is only a minor restriction: some latencies in the range of several milliseconds, occur when the graphics board is initialised for the first time. Later on, switching to and from graphics or even restarting the X server does not produce any more latencies. Since the initialisation of the graphics board can be done at boot time before the real-time critical application is started, this restriction is normally not significant." For more details visit www.osadl.org.

Azingo Mobile 2.0: A complete touch- and Web-enabled OS

Open mobile OS company Azingo has announced Azingo Mobile 2.0, a Linux-based platform that includes the Azingo Browser, Azingo Web Runtime, Azingo Application Suite, and Azingo Active Homescreen. Azingo Mobile 2.0 offers a comprehensive UI toolkit enabling a full touch user experience and Web widgets that claim to leverage device-specific services like telephony, messaging, multimedia and location-based services through the Azingo Web Runtime.

Azingo Mobile 2.0 includes all of the software, development tools, documentation and training required to design and commercialise new mobile phone products. This new platform is based on the LiMo Foundation R1 Reference Implementation. Manufacturers licensing Azingo Mobile 2.0 will receive the Azingo Browser, Azingo Web Runtime, Azingo Active Homescreen and the Azingo Application Suite.

Azingo's Active Homescreen radically extends the capabilities of a conventional phone homescreen by allowing users to add and organise pertinent, real-time information from the Internet and their phone, for fast, simple access. The Azingo Active Homescreen was designed to mimic the familiar computer desktop experience through features such as a wider, scrolling homescreen area, drag and drop, shortcuts, folders, and widgets. Users also have one-touch access to content on their handset, including native, Web, or Flash Lite applications, contacts, photos, music, videos, and messages.

The Azingo Applications Suite is available for all Linux-based mobile platforms. The Suite includes the following applications: Azingo Mobile Entertainment, Azingo Mobile Productivity, Azingo Mobile Communications, and Azingo Mobile System Applications.



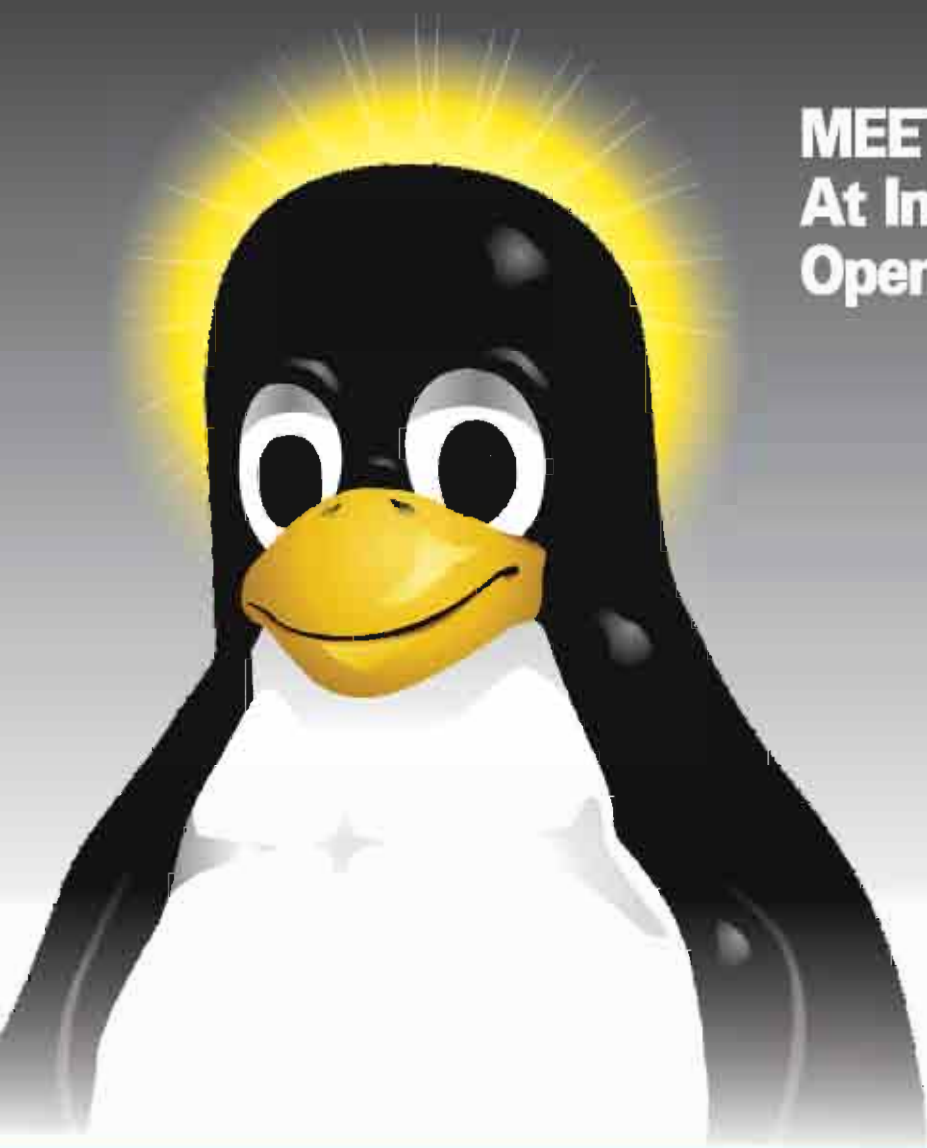
VMware View Open Client goes LGPL

VMware has announced the release of an open source client for the virtual desktop infrastructure called VMware View Open Client. VMware View supposedly enables IT organisations to safely host user desktops in the data centres, while letting users access their personalised desktop environments from almost any device, at any time. Now, the virtualisation vendor is providing VMware View Open Client for partners, enabling them to use VMware View source code to

optimise their products to deliver rich, personalised virtual desktops.

VMware View Open Client is available under the GNU Lesser General Public License version 2.1 (LGPL v 2.1) and is accessible from code.google.com/p/vmware-view-open-client. Some of the features included in this release support secure tunnelling using SSL, two factor authentication with RSA SecurID, Novell SUSE Linux Enterprise Thin Client Add-On RPM package and a full command-line interface.

**Best time
For India's IT
to grow?
RIGHT NOW.**



MEET THE MESSIAH!
At India's leading
Open Source Event

SEMINARS & WORKSHOPS



FOSS Enabled Software Development

The Promise: Develop software better, faster and cheaper

12th March (10 am to 6 pm)



CXOSummit: Virtualization, SaaS and Cloud Computing Using FOSS

The promise:

13th March (10 am to 1:30 pm)



FOSS Powered Mobile Phones & Devices

The promise: Leverage the power of FOSS for developing 'mobile' software

13th March (10 am to 1:30 pm)



FOSS For Everyone

The promise: How does FOSS benefit all of us?

14th March (10 am to 6 pm)



IT Management Simplified

The promise: IT Management: Better, simpler and cheaper

12th March (10 am to 6 pm)



FOSS enabled Education

The promise: Deploy best teaching tools yet cut costs

12th March (10 am to 1:30 pm)

TYPES OF PASSES

Silver

(Free Entry) for first 500 registrants!

• Allows entry on all 3 days • Allows access to seminars: (1) FOSS Enabled Software Development, (2) FOSS Powered Mobile Phones & Devices, and (3) FOSS For Everyone • Allows access to the Expo • Allows access to additional work-shops, seminars, and BoFs depending on their admission criteria.

Gold

• Allows entry on all 3 days • Allows access to seminars: (1) FOSS Enabled Software Development, (2) FOSS Powered Mobile Phones & Devices, and (3) FOSS For Everyone • Allows access to the Expo • Allows access to additional work-shops, seminars, and BoFs depending on their admission criteria PLUS

• Preferential seating in all seminars • Lunch, tea & coffee at the 'Networking Zone' • Certificate for attendance is provided (couriered after 2-3 weeks)

Platinum

• Allows entry on all 3 days • Allows access to seminars: (1) FOSS Enabled Software Development, (2) FOSS Powered Mobile Phones & Devices, and (3) FOSS For Everyone • Allows access to the Expo • Allows access to additional work-shops, seminars, and BoFs depending on their admission criteria • Lunch, tea & coffee at the 'Networking Zone' • Certificate for attendance is provided (couriered after 2-3 weeks) PLUS

• Reserved seating in all seminars

• Access to these seminars too: (1) IT Management Simplified and (2) CXOSummit: Virtualization, Cloud Computing & SaaS Using FOSS

For Free registration visit : <http://www.osidays.com/users/delegate09.php>

Gold Sponsors



Media Partners



Event Partner



Powered By





KNOW HOW



Q I am new to the world of Linux. I have used it on my desktop and am comfortable doing things with the graphical interface. However, a few days back I needed to rename a few files and folders on a remote computer from the command line. Please advice me on how to go about this.

—Sneha, Kolkata

To rename files and folder in Linux, you can use the `mv` command. For example, to rename a file called `test.txt` to `testnew.txt`, you can run the following command:

```
mv test.sh testnew.sh
```

Similarly, you can do this for folder/directory names too. You can even use wild cards with commands as well, if you need to rename multiple files/directories.

Q Recently, in an interview, I was asked if there was any tool to get the names and list of packages installed on Debian, Red Hat, Mandrake and SUSE. Is there any such specific tool that is used for all these distros, or are they different?

—Deeksha Chaudhary,
by e-mail

You can use the command `dpkg -l` for Debian and Ubuntu. For Red Hat, Fedora, Mandriva and SUSE, you can use `rpm -qa`.

Q I was trying to schedule some scripts to run automatically. Can you please help me understand the different settings that need to be done to set up a scheduler in Linux.

—K. Singh, Patiala

Scheduling a task/job is done using a utility called *cron*, which makes tasks automatically run in the background at regular intervals. To manage *cron* jobs, there is *crontab*, a file that contains the schedule of *cron* entries to be run at specified times. To set up a scheduler, you need to make entries in this. Executing *crontab -e* opens the file in editable mode so that you can enter the details and save it. The *crontab* syntax has five fields as mentioned below.

```
* * * * * command
```

Here, you can replace the first asterisk to enter the day of week—0 to 6, where 0 is Sunday. You can enter the month (1-12) in place of the second asterisk. Replace the third asterisk with the day of month (1-31). The fourth one is for the hour (0-23), while the fifth one is to enter the minutes (0-59).

For example, if you need to run a script daily at 5:30 p m, then the entry will as shown below:

```
30 17 * * * sh /home/myuser/
scripttorun.sh >/dev/null 2>&1
```

If the last part “>/dev/null 2>&1” is omitted, then by default, *cron* will send an e-mail to the user account after executing the *cronjob*.

Q I have an old Fedora installation on my system. I have just upgraded my RAM from 256

to 512 MB. My swap partition is also of 512 MB. Is there any way by which I can increase my swap space to 1 GB without formatting and reinstalling my OS.

—Amit Nandam, Gaya

Sure you can increase the size of swap without reinstalling the OS. This can be done by either adding a new swap partition or creating a swap file, instead. Here, I will discuss the steps to add a new swap file to increase the swap size as this is possible even if you do not have free unpartitioned space in your hard disk for creating new partitions, but have free space on one of your partitions.

To make a swap file of 1 GB, multiply 1024 MB and 1024 to get a block size. Now, at a shell prompt, as the root user, type the following command with the count being equal to the desired block size:

```
dd if=/dev/zero of=/swapfile bs=1024
count=1048576
```

Now set up the swap file:


```
mkswap /swapfile
```

Enable the swap file after creating it by using the following command:

```
swapon /swapfile
```

Add the following entry in the the */etc/fstab* file to make the system activate this file as swap while booting the system:

```
/swapfile swap swap defaults 0 0
```

This will enable your 1 GB swap space, every time your system boots. You can check the size of swap by using the *free* command or by *cat /proc/swap*. 

17th Convergence India 2009

International Exhibition & Conference

Pragati Maidan, New Delhi, India

18-20 March 2009

*South Asia's largest **ICT** Event*
*Digital Convergence changing the **ICT** Landscape*



17th Convergence India 2009

Conference Participation Charges

3 day sessions	Rs 10,000 / US\$ 225 Per Delegate
2 day sessions	Rs 7,000 / US\$ 175 Per Delegate
1 day session	Rs 4,000 / US\$ 100 Per Delegate



Certified by



Department of Telecommunications
Ministry of Communications & Information Technology
Government of India

Supported by



Department of Information Technology
Ministry of Communications & Information Technology
Government of India



Ministry of Information & Broadcasting
Government of India



Supporting Journal



Organised by



Exhibitions India Pvt. Ltd.
(An ISO 9001:2000 Certified Company)

For Conference Registration, please contact:

Shveta Sethi, shvetas@eigroup.in; Rahul Torani, rahult@eigroup.in; Divya Tiwari, divyat@eigroup.in; tabisha@eigroup.in Tel: +91 11 4279 5000

217-B, (2nd Floor) Okhla Industrial Estate, Phase III, New Delhi 110 020, India Tel: +91 11 4279 5000 Fax: +91 11 4279 5098/99

Bunny Sidhu, Vice President, (M) +91 98733 43925 bunnys@eigroup.in /

Sambit Mund, Group Manager, (M) +91 93126 55071; sambitm@eigroup.in

Branches: Bangalore, Chennai, Hyderabad, Mumbai, Ahmedabad, California

www.convergenceindia.org

Gee, I Like Your Desktop!



The newly released version KDE 4.2 stands out because it offers a fantastic desktop experience.

While your system boots the live CD (built on top of an openSUSE 11.1 base), you will be looking at that same old dull green boot screen of openSUSE. Wondering why I am picking on openSUSE's default theme again this month? Well, that's because the KDE 4.2 Live CD, which is bundled with this issue of *LFY*, is based on it. Thankfully, once the desktop loads up, you are greeted by the default look and feel of the official desktop release. That's the desktop shell dubbed Plasma by the way, whose job it is to let you organise your desktop pretty much the way you want it.

Oxygen: Breathe in, breathe out!

Yes, that default look and feel that KDE 4.2 comes with is thanks to something called Oxygen—the theme, the window borders

and those beautiful icons. The noticeable change this time round is the desktop panel, which is now a shade of blue—much prettier, I must add.

The panel

The panel has the usual stuff—the KDE (Kickoff) menu, the *Show desktop* utility, the pager (or the workspace switcher), the task manager, device notifier and system tray, followed by the clock. I don't remember if that's the exact order of things, by default, but that's how I like it. If you're on a laptop, the panel should also have the battery monitor next to the system tray.

Apart from a new icon for the device notifier, the utility that has attracted a significant amount of attention from developers is the system tray—the thing that generally holds the KMix (sound

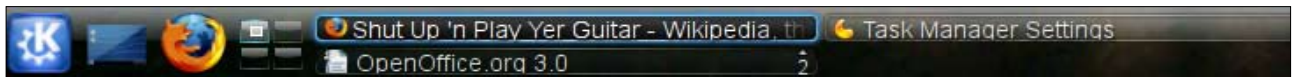


Figure 1: Tasks in Task Manager organised in two rows

control), Klipper (clipboard), and perhaps a software update notifier too. The tray now looks like, well...a tray, due to it being confined by a boundary. Also, you can right click on its edge to get a hold of the *System Tray Settings* menu, from where you can select icons of the running apps you want to be auto hidden. After you've made your selections, the size of the system tray will get shortened, and the left side of the tray will display an arrow pointing left, indicating you can expand it in that direction. This option comes in handy if you run a lot of apps that place a control icon inside the system tray, thus expanding the tray to occupy the precious panel space.

The task manager has also got a bit of a facelift—it can now again group applications based on program names, and placing the tasks on multiple rows is also possible just like in KDE3. All this can be configured from the *Task Manager Settings* menu by right clicking on the panel. Although, I have no complaints about how grouping works, when you organise tasks in rows, the default teaming makes it look a bit odd—as if someone has shrunk the tasks forcefully (Figure 1). However, I'm sure this would be addressed in a bug fix release soon, as it looks like a theme issue. Also, by default, while tasks are sorted in alphabetical order (and not in the order programs are launched) I've enabled it to let me sort stuff manually. This enables me to drag and reorganise my programs as I wish.

In addition to all this, more options have been added to the general settings of the overall panel. Things like how to increase the height, position and screen-edge can be much more easily achieved now, unlike in KDE 4.1, where it was much harder to guess how to go about things.

Coming back to those who're on the move, right click on the



Figure 2: My personal desktop with Folder View, Comic Strip, Notes and Picture Frame widgets

Battery Monitor widget (should be somewhere near your clock), and you'll see configurable options galore. This is thanks to the integration of the PowerDevil utility, another addition in this release. It offers various pre-configured 'Power Profiles'—viz. performance, powersave, aggressive powersave, presentation, etc.—and lets you fine tune all the profiles as per your liking. Overall, although it's pretty easy to use and understand the options, KDE4 seems to drain out a lot of battery power compared to KDE3 and GNOME.

As for the default KDE4 menu, Kickoff, it hasn't got any visually noticeable feature additions, apart from the border colour, which is now black, to gel with the rest of the Oxygen theme.

Workspace

As for the other part of the desktop, which is the main workspace, things have again been refined a lot with respect to the widgets, and their numbers have also increased

considerably. Figure 2 shows what my personal desktop screen looks like, while Figure 3 shows my work laptop. As you can see on my personal desktop, I have the Folder View widget, a Calvin and Hobbs comic strip, plus a few pictures of my family members.

The Folder View widget, as you know, was introduced in KDE 4.1. You configure it to point towards the contents of the folder you access frequently, right on your desktop—so there's no need to use the file manager to hunt down *that* folder every time you need something; however deep inside the filesystem it is located, you can see its contents right on your desktop. You can also have multiple Folder Views, say *~/Documents* and *~/Pictures* folders, for easy access.

Traditionally, the contents of the *~/Desktop* folder are displayed as icons on your desktop. Well, by default, the Folder View widget displays the contents of this folder. In fact, if you like, in KDE 4.2 you can even make your desktop imitate



Figure 3: My work laptop with the Folder View, Notes, Picture Frame and RSSNow widgets

the traditional versions, with icons and files all over the place. You can set it by accessing the *Appearance Settings* option by right clicking on your desktop and then changing your desktop activity type from 'Desktop' to 'Folder View'. But I don't think it's a good idea, especially if you are someone like me who keeps downloading random stuff from the Web, and storing it on the ~/Desktop folder, turning the desktop into a huge pool of icons. That's why I would rather use Folder View as a widget than use it as *my desktop*. The reason being that I can set its size to what I want and use the rest of the desktop to put other useful widgets, without worrying about my ~/Desktop folder becoming a junkyard of trash downloads from all over the Web.

This brings me to the other useful widgets I use. The Calvin and Hobbs is courtesy the Comic Strip widget, and it works provided you're connected to the Internet. When you launch it, you first need to set it up. You can pull comic streams from a wide range of streams hosted at KDE-Files.org—so whether you're into Garfield, Dilbert, or anybody else, you are free to choose from a list that's more than a handful. This is a healthy addition considering v4.1.x only provided me with an option of a few.

The pictures are courtesy the Picture Frame widget. Here, you can simply drag and drop pictures and the widget adjusts its size according to the dimensions (landscape or portrait) of the source image. Don't worry, the widget doesn't depend on the resolution of the source image—it's intelligent enough to give you a good-sized picture frame, and you are free to increase or decrease it as per your taste. This is a very useful addition for me, as earlier I used to remix an image with lots of elements (and faces) to create wallpaper. Now, I can simply select an image of some scenery for a wall paper, while pictures of people go into frames. :-)

For my work desktop, I like to track a few news sites and the RSSNow widget lets me do exactly that. Each feed automatically scrolls horizontally to show me the current news—and of course, it gives me the option to manually skim through them. When I chance upon something interesting, I click, and the default browser loads the Web page with the complete story.

Another handy widget is pastebin. If you hang out in IRC channels, I don't need to explain what pastebins are. You probably point your browser to a pastebin website, upload the information and then obtain the URL to share with others. Instead, the widget connects



Figure 4: Lancelot launcher

to *pastebin.ca* for you, so all you need to do is drag and drop text or images here, and it gives you the URL that you can post in the IRC channel you're logged in to. Makes life a lot easier, right?

Other widgets that you may find handy are the age-old binary clock, Blue Marble (a 3D model of the earth that's rendered thanks to the Marble application—a Google Earth-like tool), Calculator, Dictionary, Eyes, Fifteen Puzzle, LCD Weather Station (to keep an eye on the current weather of your city), Luna (to check the current phase of the moon), Twitter Microblogging client, World Clock, various system monitors (to monitor your hard disk, CPU temperature, network traffic and other hardware information), besides a lot more that don't really seem interesting enough to me. I also hope some of these widgets get the attention of the artists teams to make them look more appealing, like the system monitors that take up too much screen space and look too dull.

Before ending my rant on the widgets, allow me to draw your attention to the Lancelot Launcher menu in Figure 4. Although, technically it's more or less similar to Kickoff, I like the way things are organised here, besides the fact that it looks more appealing. After customising a few of its settings, I've finally switched to Lancelot as my default launcher menu. Although, I hardly use even *this*: which brings me to...



Figure 5: KRunner

KRunner

This is the 'Run' command that you activate by pressing **Alt+F2** from the keyboard. Although it's been available for a year now, things have been more aesthetically refined in this release. You can make it work in either of two modes—command-based like we've been used to since ages, or now even task-based. By default, you can key in the commands to launch an application—as soon as you start typing, it starts filtering from the names of the applications available. For example, take a look at Figure 5—as soon as I type 'Kon', it filters from all the commands for apps that contain the letters 'kon'. I can now either key in a few more letters to fine-tune the filtering further, or use **Tab** to select the program I need to launch.

You can also use KRunner for a lot of other purposes, viz., as a calculator, to find documents, search by tags, or even visit website URLs you directly key in. This extended functionality is courtesy several plug-ins that power its back-end. A few examples are spell checking, browser history, recent documents, etc. You can check out the full listing of plug-ins by accessing its configuration dialogue. In fact, this is where you can switch from command-based to task-based mode by activating it from the 'User Interface' tab. As you can see in Figure 6, now you can find applications by their task, rather than command—for example, I typed 'write' and it shows me all the applications that can help me write something. However, since I'm too used to the command-based mode, I

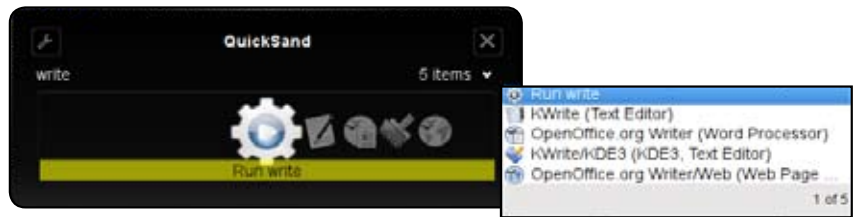


Figure 6: QuickSand—KRunner in task mode



Figure 7: The Alt+Tab effect

found the task-oriented method kind of difficult to use.

Talking about the command mode, thanks to the back-end calculator plug-in, I can now see the outcome of simple math problems from KRunner itself, without launching the calculator application separately. For example, type '2134*134=', excluding the single quotes. Did you see '285554' right away? Pretty cool, eh?

Overall, KRunner is not just a regular Run dialog any more—it's turning out to be a pretty powerful application in itself.

Kwin

Kwin is basically the window manager—something that acts as a container for the apps we run on our desktop. Even in KDE 4.1, we saw some pretty cool compositing and desktop effects features added

to Kwin. This release has added even more plug-ins for effects, and the ones that were already available, have been fine tuned—refer to Figure 7 for an improved Alt+Tab effect, when too many windows are open; note the horizontal viewer at the top. You can activate it from the 'Desktop' settings in the Personal Settings app (the replacement for KControl from the KDE3 branch).

I won't go on and on about the desktop effects it offers; you should try it out yourself to experience its niceties. Maybe you'll find a lot less polish with respect to some features that Compiz Fusion also offers, but things aren't that bad either. In fact, looking at the features made available during the last six months, I won't be surprised if it catches up with Compiz by the time KDE 4.3 comes out.

As for me, although I usually keep

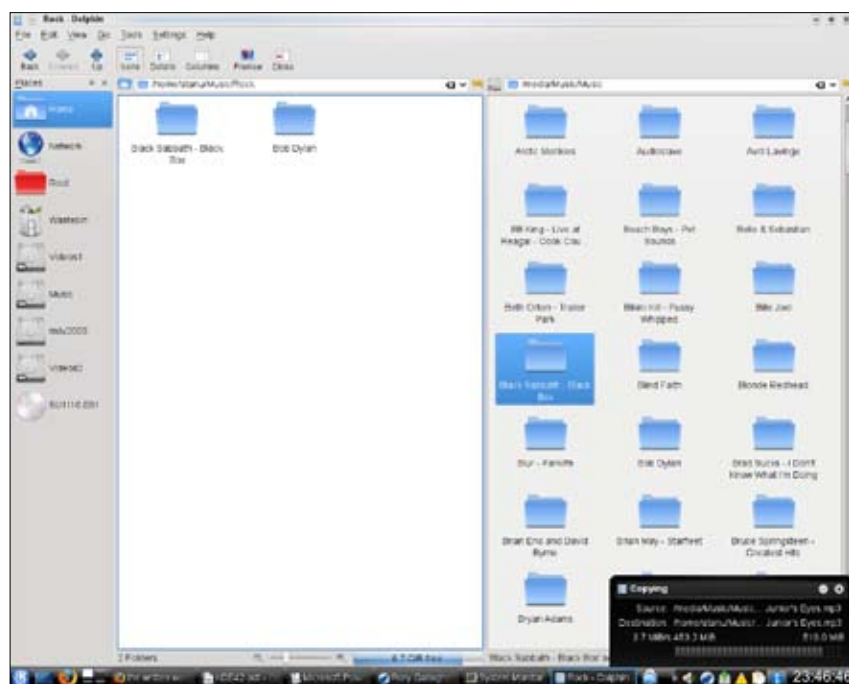


Figure 8: Split view in Dolphin file manager

these effects disabled as I find them distracting. I don't mind playing around with them when I need to kill time, or show off in front of those Winduhs users. And for that purpose, I'd rather prefer a native window manager to take care of the effects rather than using a third-party tool, which most of the time asks me to log out and log back in to activate/disable the effects. Although, I've got to admit, these effects work much better on Intel's graphics, rather than the proprietary drivers that ATI and Nvidia depend on—mostly due to bugs in the drivers.

A working man's place

Well, at the end of the day, with all these great features, you still need applications to survive. KDE 4.2, with its accompanying applications, doesn't disappoint here either. Whether it's a personal information manager, Internet apps, image and document viewers, media players, desktop administration tools, or other assorted utilities, it's got you more or less covered everywhere. And once the ambitious KOffice 2 goes stable (currently the CD includes the beta version), KDE4 could become an all-in-one desktop

powerhouse.

Although, we from the free software world have grown too used to OpenOffice.org for an office suite, KOffice 2 indeed offers a pretty useful alternative, which consumes a lot less memory compared to the former, besides introducing a pretty innovative user interface and set of features. The same is true if you compare Konqueror as a Web browser with Firefox. However, both KOffice and Konqueror have some catching up to do before they can pose a threat to the dominance of OpenOffice.org and Firefox, respectively.

In fact, now that KDE has got a default file manager called Dolphin, I hope the Konqueror developers concentrate more on its Web page rendering capabilities to make it compatible with the websites around the globe. One good option is to dump the KHTML engine for Webkit. Although Webkit is available as an optional rendering engine, the pages still appear ad hoc—which leads me to wonder what the issue might be, because I hear Google Chrome uses the same engine.

Anyway, for now, since we have OpenOffice.org and Firefox to take

care of our requirements, let's look at some of the other programs that help us in our day-to-day work.

Dolphin

The most important tool is obviously the file manager. This is what you essentially use to browse the gigabytes of data stored in your hard drive, which you can't do without. Dolphin concentrates on taking care of exactly this, unlike Konqueror in KDE3, which sort of posed as an all-in-one tool for multiple purposes.

The Dolphin interface is simple, with no frills at all. Okay, that information pane on the right looks slightly weird, you say? Well, it's there to display information about the file you select. But, if you want to get rid of it, simply hit F11 and it's gone. To tell you the truth, I don't like it either, and would rather have tool tips provide me the information when I hover my mouse over a file. That's doable too; you can configure all this from the *Configure Dolphin* option under the *Settings* menu.

A new addition in this release is the zoom slider at the bottom right corner of Dolphin—you can use it to increase the size of icons in the file manager. However, if you remember Dolphin from KDE 4.1, it used to display the free disk space in the same location, which is a nice way to keep an eye on your disk activities. What made it disappear this time? You got me there! But, again, you can enable it to appear beside the zoom slider from the *Configure Dolphin* option.

And here's two handy shortcuts for you, just in case you aren't aware of them already: if you don't like the breadcrumbs-based navigation bar, press Ctrl+L and you get the traditional input field to enter a location. If you want to filter a file from a directory listing hundreds of files, press Ctrl+I to get the filter input box at the bottom of the window. And if you want these features permanently available, visit *Configure Dolphin* again.

Dolphin does offer some

advanced features too. One of my favourites is the 'Split' view option to divide the window into two parts—activate by F3 or *View→Split*. This option (Figure 8) comes in very handy when you want to copy/move files between two locations. Simply drag and drop the file/folder from one side to the other. By the way, the copying dialogue box has a nice notification window now, that pops out of the system tray, which even gives you the option to pause a transfer—much better than a separate floating window, I'd say. This dialogue box is not specific to Dolphin, but ubiquitous for copy/move/download/upload activities across all KDE apps.

Another good addition is the preview option in the toolbar. Now, I don't need previews for regular files, but only in a directory of images. When I go inside the images folder, I hit preview and use the zoom slider to increase the size of the thumbnails; Dolphin remembers this setting. So, next time I launch Dolphin, although the rest of the folders appear as they were, inside the Images folder, the preview option is still active, and the thumbnail sizes are still bigger. Neat, eh?

Other assorted apps

Some of the other KDE-specific tools I have to depend on heavily for my daily grind are personal information manager Kontact, Gwenview image viewer and Okular document viewer. All of them are full-fledged apps with pretty advanced capabilities, and deserve a few pages to cover the features they each offer.

Kontact is something I depend a lot on at work. Apart from the Kmail client and an address book, it gives me an RSS feed reader, a to-do list, a calendar, time tracker, and something called pop-up notes. I absolutely love the addition of the 'Fancy with Clickable Status' theme in Kmail, which finally made me switch to a vertical message preview pane using a three-column layout in

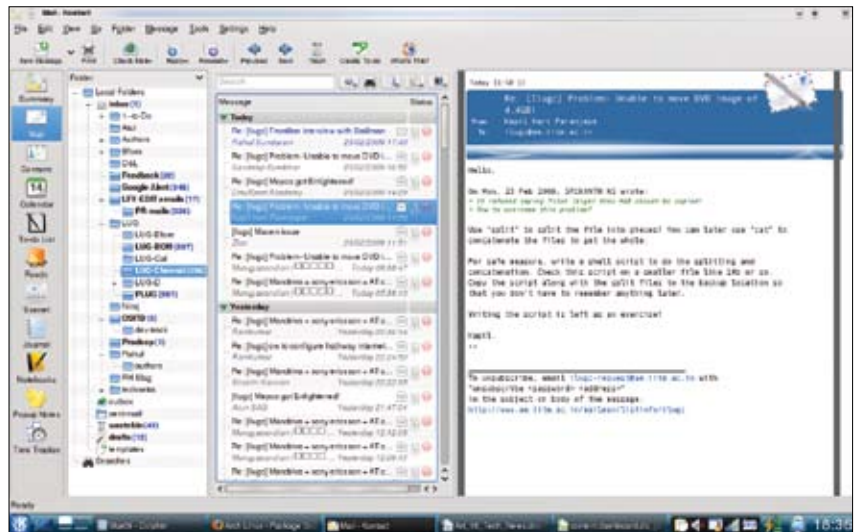


Figure 9: The 'Fancy with Clickable Status' theme in Kmail

my wide screen laptop (Figure 9).

My image viewing needs are fulfilled by Gwenview, which offers me some handy image editing features, apart from working as a powerful image viewer. Okular, on the other hand, takes care of displaying a range of document formats. For your digital camera requirements, there's Digikam, although the Qt4 version is still in its beta stage.

Kget is available for normal downloads. The interface is very simple to use, while it also lets me unleash a lot of advanced features when I need them. It even supports torrents and metlinks for downloads. Although, for torrents, I like sticking to KTorrent, which is a full blown torrent client with some fantastic features.


KDE 4.2 comes with JuK for an audio player and music collection manager, which does what it's supposed to do, pretty well. However, most of us are too used to Amarok anyway, which has also finally released its Qt4-based stable v2.0, with some intuitive features. Dragon Player is the default media player in the KDE4 series, which is more like its GNOME counterpart Totem, that does its job with no frills attached. Those who need more functionality anyway have the MPlayer-based SMPlayer, which I recommend

that you check out just in case you haven't already.

Besides these, there're lots of other tools in the accompanying KDE 4.2 live CD; it's just that I don't have enough space to write about them here. Explore for yourself—there are the educational tools, the highlight being the Google Earth alternative, Marble, and various other utility programs.

Oh wait, before I end: don't forget to check out the much-improved System Monitor application. The highlight of this release is the *System Load* tab. The interface has finally been made similar to its GNOME counterpart, which always had a much better UI.

Bottom line

Well, KDE 4.2 is not perfect! It still tends to crash on me occasionally for no obvious reason, but things are getting there. And even if it's still not as feature complete as KDE 3.5.10, with this release things have indeed gotten a LOT better. So, c'mon, don't be so uptight and give it a go! Or, do you really want to wait another six months for version 4.3? Well? **END** 

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

STOP WASTING CDs

Install Linux Straight from an ISO

You download the brand new Debian 5 (when it's released) after waiting for so many months, and discover you don't have a single blank DVD to burn the ISO image! Why worry, when there's a simple way out!



*G*NU/Linux comes in many different flavours, apart from the fact that each individual distro has a new release almost every six months, if not less. I have a habit of trying out every new version the moment it comes out, and I'm sure many of you do too.

Now, let's assume you have downloaded a new version of a distro and are in the mood to try it out right away. It's past midnight and you realise that you've run out of blank CDs/DVDs. So you will have to wait till the morning when the shops open, to be able to burn the distro image in order to install it. I'm sure a lot of us often face this problem. In this article I'll share a simple trick by which you can install the new distro without burning it to a CD/DVD. The only requirement

is that you should have a pre-installed GNU/Linux system—which you already have, otherwise where did you download the ISO image from?

All Linux installers use two files to boot a computer: a kernel and an initial root filesystem -- also known as the RAM disk or *initrd* image. This *initrd* image contains a set of executables and drivers that are needed to mount the real root filesystem. When the real root filesystem mounts, the *initrd* is unmounted and its memory is freed. These two files are named differently in different distros—refer to Table 1 for their names.

The first thing you need to do is place the ISO image(s) inside a directory. Some installers are not able to read the ISO images if they are placed inside a directory. So, just to be on the safe side, place them in the root of the file system. The partition on the hard disk holding the ISO files must be formatted with the ext2, ext3 or vfat file system.

In our example, let's go ahead and do it with an old Fedora 9 ISO image. Follow these steps to begin with:

```
# mkdir /fedora
```




```
# cp /home/sandeep/Fedora-9-i386-DVD.iso /fedora/fedora9.iso
```

Now extract the kernel and initrd files from the ISO image and place them in the same directory in which you placed the ISO. You can use File Roller, the archive manager for GNOME, to extract the files. Just right click on the ISO and select *Open with File Roller*. It displays the contents of the ISO image. Then navigate to the *isolinux* directory—in Fedora 9 these two files are placed inside the *isolinux* directory; it's often different for other distros, so please refer to Table 1 for the paths. Select the kernel and initrd files, and extract them to the location where your ISO image exists.

The second method is to mount the ISO image and extract the files. Run the following commands to do this:

```
# mount -o loop /fedora/fedora9.iso /media/iso
# cd /media/iso/isolinux
# cp vmlinuz initrd.img /fedora/
```

I have mounted the ISO image without providing the *-t iso9660* option (to specify the type of media as an ISO filesystem). It worked for me. If the above mount command doesn't work, do add this option along with the rest of the *mount* command above.

 **Note:** Fedora 10 has introduced a change in the Anaconda installer. So, in addition to the *vmlinuz* and *initrd.gz* files, you will also need to copy the *images/install.img* file, create a directory called */fedora/images*, and place the *install.img* file there.

Now, it's time to edit the */boot/grub/menu.lst* file on the system I'm currently using—Ubuntu 8.10. Note that this is the location of the Grub menu in almost all distros, except for Fedora/Red Hat, where it's called */boot/grub/grub.conf*. Append the following entry there:

```
title Install Linux
root (hdX,Y)
kernel /distro/Linux_kernel
initrd /distro/Ram_disk
```

In this case...

- 'title' is the name you want to display in your GRUB menu
- 'root' is the hard disk partition that contains the ISO image
- 'kernel' is the Linux kernel
- 'initrd' is the initial RAM disk image

Likewise, the *menu.lst* entry for the ISO file looks like what's shown below:

```
title Install Fedora 9
root (hd4,0)
kernel /fedora/vmlinuz
initrd /fedora/initrd.img
```

Names of kernel and RAM disk images in some popular distros

Linux OS	Kernel path	Ram disk path
Fedora	/isolinux/vmlinuz	/isolinux/initrd.img
openSUSE	/boot/i386/loader/linux	/boot/i386/loader/initrd
Mandriva	/i586/isolinux/alt0/vmlinuz	/i586/isolinux/alt0/all.rdz
Ubuntu	/casper/vmlinuz	/casper/initrd.gz
Debian	/install.386/vmlinuz	/install.386/initrd.gz
RHEL5/CentOS5	/isolinux/vmlinuz	/isolinux/initrd.img

Table 1

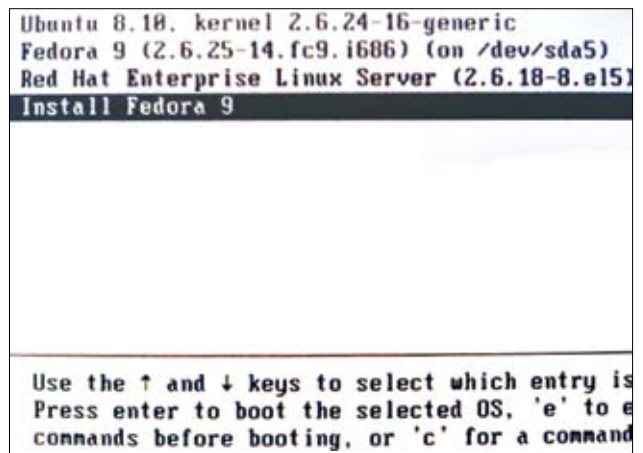


Figure 1: The Install Fedora 9 Grub menu entry

Now you are ready to install your new Linux distro directly from the hard disk without the need for a CD/DVD drive. Reboot your system and select the 'Install Fedora 9' entry from your GRUB menu.

Figure 1 shows what the GRUB menu looks like after rebooting my system.

Obviously, I selected the 'Install Fedora 9' entry and it has started booting my system with the help of *vmlinuz* and *initrd.img* files. The set-up prompts me to choose a language and keyboard layout. Then it prompts me to select the 'Installation Method' as shown in Figure 2.

In this screen you need to select the 'Hard drive' option and proceed to the next screen. Here, you have to select the appropriate partition and the directory where the installation image exists. In my system, the installation image exists in the */fedora* directory of */dev/sda5* partition. This is shown in Figure 3.

...continued on page 27



Figure 2: Select the hard drive for 'Installation Method'



Let's have some fun with Linux commands.

Many of us who love to work on Linux enjoy the privilege of using a plethora of commands and tools. Here is our effort to introduce you to a few very simple- to-use, yet enormously effective commands. The intended audience may belong to all classes of Linux users and the only requirement is to have a basic acquaintance with Linux. Our article deals with bash shell and Linux version Fedora 9, kernel 2.6.25.

- a)** Often, commands on the console may span many lines, and encountering a type mistake at the beginning of the command would require you to use the slow way of punching the right/left arrow keys to traverse in the command string.

Remedy: Try **Ctrl+E** to move to the end of the command string and **Ctrl+A** to reach *Start*. It's the fastest way to edit a Linux command line. To delete a word in the command string, use **Ctrl+W**.

- b)** Another wonder of a simple shell variable is **!\$**. Let's say you have to create a directory, go into it and then rename it. So the flow of commands would be:

```
$ mkdir your_dir
$ mv your_dir my_dir
$ cd my_dir
```

Remedy: Well, Linux has a shorter and quicker way:

```
$ mkdir your_dir
$ mv !$ my_dir
$ cd !$
```

!\$ points to the last *string* in the command string. This is useful in various scenarios where the last word of command string is to be used in subsequent commands (almost with all Linux commands like *vi*, *tar*, *gzip*, etc).

- c)** Do you want to know what an *ls* or a *date* command does internally? Just run the following code to get to know the basic block of any Linux command:

```
$ strace -c /usr/bin/ls
```

strace is a system call monitor command and provides information about system calls made by an application, including the call arguments and return value.

- d) What if you want to create a chain of directories and sub-directories, something like `/tmp/our/your/mine`?

Remedy: Try this:

```
$ mkdir -p /tmp/our/your/mine
```

- e) One very interesting way to combine some related commands is with `&&`.

```
$ cd dir_name && ls -alr && cd ..
```

- f) Now for some fun! Have you ever tried checking the vulnerability of your Linux system? Try a fork-bomb to evaluate this:


```
$ :(){&};:
```

It's actually a shell function; look closely and it's an unnamed function `:()` with the body enclosed in `{}`. The statement `'|:'` makes a call to the function itself and pipes the output to another function call—thus we are calling the function twice. `&` puts all processes in the background and hence you can't kill any process. Finally `;` completes the function definition and the last `:` initiates a call to this unnamed function. So it recursively creates

processes and eventually your system will hang. This is one of the most dangerous Linux commands and may cause your computer to crash!

Remedy: How to avoid a fork bomb? Of course, by limiting the process limit; you need to edit `/etc/security/limits.conf`. Edit the variable `nproc` to `user_name hard nproc 100`. You require root privileges to modify this file.

- g) One more dirty way to hack into the system is through continuous reboots, resulting in the total breakdown of a Linux machine. Here's an option that you need root access for. Edit the file `/etc/inittab` and modify the line `id:5:initdefault:` to `id:6:initdefault:`. That's all! Linux specifies various user modes and 6 is intended for reboot. Hence, your machine keeps on rebooting every time it checks for the default user mode.

Remedy: Modify your *Grub* configuration (the Linux bootloader) and boot in single user mode. Edit the file `/etc/inittab` and change the default user level to 5. I hope you'll have some fun trying out these commands, and that they bring you closer to Linux. Please do share your feedback and comments. **END** 

By: Anshu Bhola and Vishal Kanaujia

The authors are with Hewlett Packard, India and work in the compilers and tools development group.

...continued from page 25: Stop Wasting CDs




Figure 3: Select the partition and the sub-directory where the ISO image resides

After this, it picks up the Anaconda installer of Fedora 9 (or any other installer, as in your case) from the prescribed location, and proceeds with the regular installation procedure just like you'd get if you were installing from a bootable optical media. Follow the steps as you would to install the distro. Figure 4 shows the package installation in action. After that's done, reboot and you'll be able to use your newly installed operating system.

Easy enough, right? So, I hope you'll start using this simple



Figure 4: Fedora 9 installation in action

trick to install the newly released GNU/Linux distros and stop worrying about whether you have the required blank optical media. And the additional environmental benefit is less use of non-biodegradable plastic materials (which is what a CD/DVD is made out of). **END** 

By: Sandeep Yadav

The author is a part of the LFY CD team and loves to run `./configure && make && make install` every now and then.



Managing

Music Efficiently

Are your audio files scattered all over your hard disk with missing meta data, leaving you with no easy way to recognise the songs? It's time you got a bit organised!

Can you imagine a life without music? It's said that music has the power to heal diseases. Music has the ability to change moods, soothe minds and lower tensions. Music is a great way to recover from stress and pain. And, thankfully, music has come of age from LPs and cassettes, to optical media. And now, with computing becoming a *de facto* standard, everyone prefers listening to their music on PCs.

PCs have become media hubs thanks to the astounding development in media formats. Ease of audio management has added to PCs becoming the centralised media for all entertainment needs. Storing music in PCs has a great advantage over other storage

resources. It saves space, lessens optical disc usage, and having a soft copy around the corner makes it easy to access it anywhere through a network or the Internet.

And today, portable media players (PMP) make it even easier to access music.

However, for everything to work impeccably and to get the most from the environment, it's necessary to maintain and keep things in order. To get the best out of our PCs and PMPs, and use the latest features, we need to manage our music collection.

Manoeuvring through music collections might be time consuming and irritating, but once you are done with it, you'll only benefit by it. There are several reasons that make

music management worth investing time in. Consider a library in which everything is haphazardly arranged! Will you be able to find what you're looking for? Similarly, I'm sure it gets on your nerves when you can't find the song you are looking for. After a few hours of tagging and organising the songs, life can become much easier.

The benefits of a properly-tagged and managed music collection are:

1. Easier and faster search
2. The ability to use advanced music features like album art, cover switcher (in iPods only), etc.
3. Identification of albums and artists

To get started with managing your music, all you need are a few supplementary tools. So here is a guide to editing and managing your music effectively in GNU/Linux.

You will require a few specialised tools and media players:

- EasyTag: An advanced music ID3 tag editor
- Picard: An online music editor from MusicBrainz
- An audio player
 - Banshee 1.4.1 for GNOME users
 - Amarok 2.0/1.4.10 for KDE

Note that you can always select the media player you like; I have chosen Amarok and Banshee because these are the most advanced media players from the Linux barracks. Since Amarok 2 is still under development, you might find some glitches while working with it. The latest version of Amarok requires KDE 4.1.3; so if you have KDE you should consider updating your DE first. Additionally, Mandriva KDE users might face some problems playing a few audio files with Amarok because Mandriva 2009 KDE includes *phonon-gstreamer*. To pass this problem, first install *phonon-xine* and then remove *phonon-gstreamer*.

Installing and adding software

Installing the software is perhaps the easiest part. You can install all the above-mentioned software using your distro's default package manager, considering that all the major distros have the latest stable packages in their repositories.

If there is no package available for your distro, you can always download the source code from the project site [see *Resources* section] and compile it. You can also take the help of the following websites to locate a binary package for your distribution:

- rpmfind.net: A one-stop shop for RPM users, which lists RPM packages provided by most of the RPM-based distros.
- getdeb.net: This website not only provides easy access to deb packages for Debian-based distros, but also gives quite a lot of information about the software you are downloading.

Note that users of Ubuntu Hardy Heron (8.04 LTS) or earlier versions will not find Banshee 1 in their default repositories. In order to get the latest version of Banshee, either install from the source code or add Banshee PPA

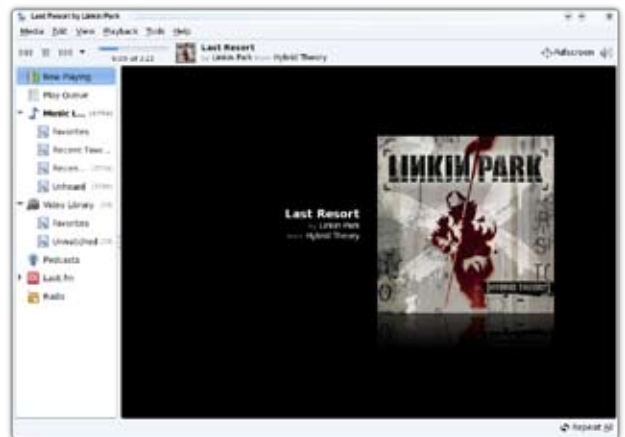


Figure 1: 'Now Playing' in Banshee



Figure 2: The default interface of Amarok 2 RC1



Figure 3: Banshee, showing the added albums

repositories—this gets you the latest stable packages of Banshee. To add this repository, click on the *System* menu on the GNOME panel, navigate to *Administration*→*Software Source*, and add the repositories listed in edge.launchpad.net/~banshee-team/+archive

Mandriva users need to enable the *testing* repositories. Navigate to *Mandriva Control Centre*→*Configure Media* and enable it.



Figure 4: Amarok displays the song playing currently

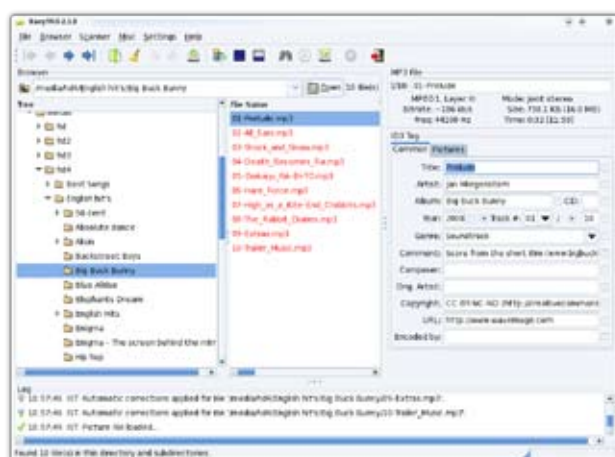


Figure 5: The EasyTag editor

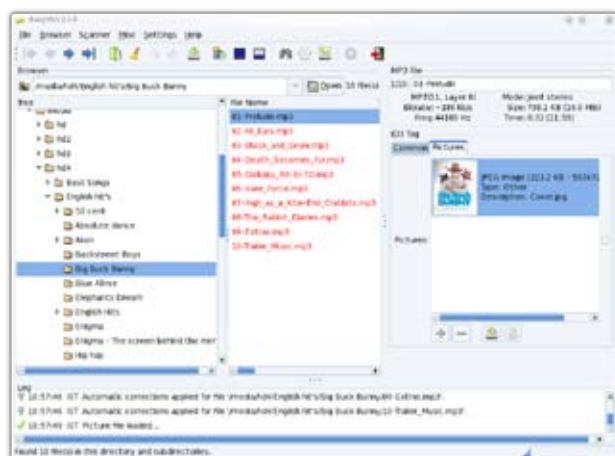


Figure 6: EasyTag's album art editor

Managing music and editing tags

The first step to managing music is to store your collection under proper directories with suitable names. Doing this will help you find or access your music without any hassle.

Whether you get your music from the Internet, from optical media, from friends or any other source, make sure you create proper folders to help you differentiate between

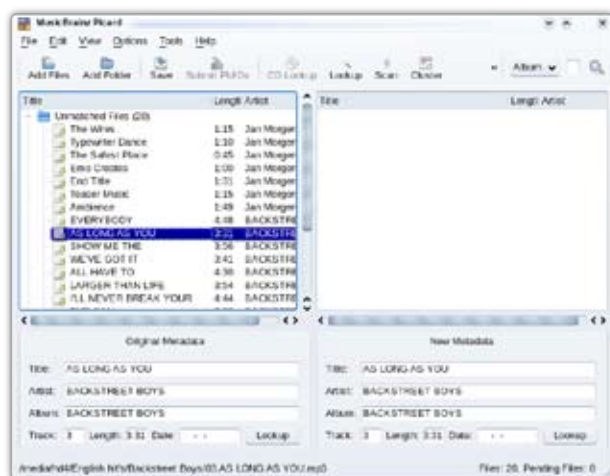


Figure 7: The Picard online tags editor from MusicBrainz

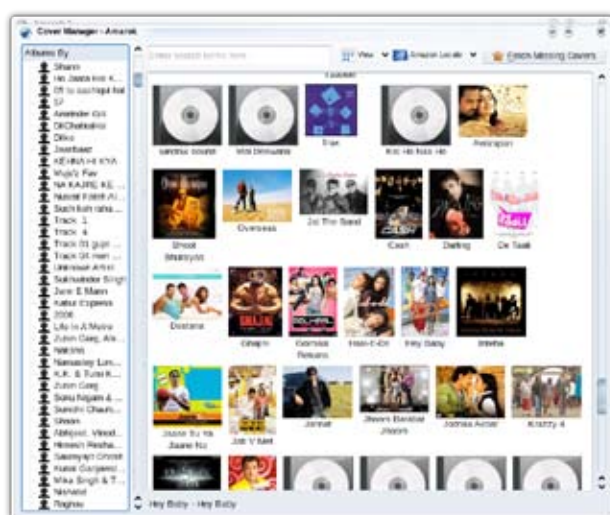


Figure 8: Amarok's cover manager

different artists and their albums. Then copy the audio files into suitable folders. Rename the audio files by track number, followed by the name of the song.

After you have stored your music systematically, it's time to move on to the next step: tagging and adding album art.

Most of the media players available for Linux come with in-built ID tag editors, but I would still recommend you use EasyTag and Picard. The editors that media players bundle, lack functionality. They allow you to edit quite a few tags, but don't allow you to edit album art. It would be great to see a full-blown ID tag editor in one of the major media players.

We will use EasyTag for general tag editing. It's a standalone tag editor with support for almost all the media formats out there—the currently supported formats are MP2, MP3, FLAC, Ogg Vorbis, Mp4/AAC, Mousepack and Monkey audio files.

EasyTag has a very intuitive interface (Figure 5). On the left-hand side you will notice the browser pane, with which you can navigate your collection. The centre portion lists all the music/audio files available in the folder that you select

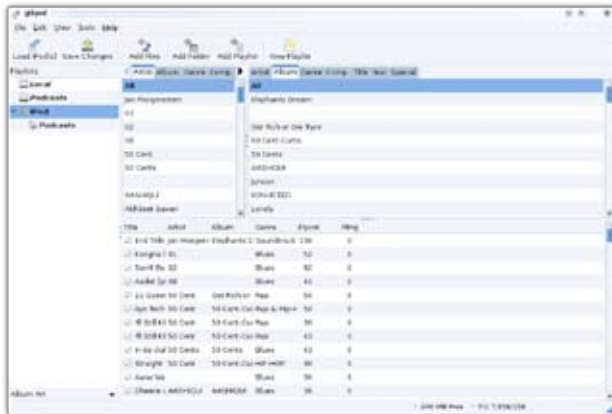


Figure 9: gtkpod in action

from the left-hand side pane, while on the right-hand side you find fields to fill/edit ID tags of the music file you select in the centre pane.

EasyTag provides an option to remove/add album art in the audio file. This is a great function for users who own portable media players with LCD screens. Users can watch and navigate through albums with the help of album art, under these circumstances.

Considering that hardware prices are continuously falling and that hard drive capacities are growing, users tend to increase their data collection. The same goes for media files. So, if you belong to that same race, it would be really hard for you to remember the ID tags for every media file. This is where Picard comes into action. Picard is an advanced ID tag editor that not only edits tags, but also uses the online MusicBrainz service to suggest ID tags for a particular track.

MusicBrainz has a voluminous collection of audio files, ranging from popular music to regional tracks. You can submit new music meta data at the website either through the software or using your Web browser.

After using Picard, I guess you are now ready with a completely 'dressed-up' music collection. However, we should add the final ingredient as well—the album art.

There are several ways to do this. The simplest is to leave it to your media player. Applications like Banshee and Amarok (Figure 8) automatically fetch album art and assign it to your album. They work magnificently with a lot of albums, but also fail when your album resembles the name of another international movie/album, in which case it may assign the wrong artwork.

To get the exact album art, you need to depend on Google/Wikipedia, or the website of the recording studio. Search online by album name for the album art/cover posted and save it to the respective album folders. Rename the file as 'Cover' or 'Album Art'. One thing to note about applications like Amarok and Banshee is that they automatically add the album art from the folder and will not fetch it from the Internet.

Owners of media players like Creative Zen, Cowon iAudio and others can easily add music by dragging and

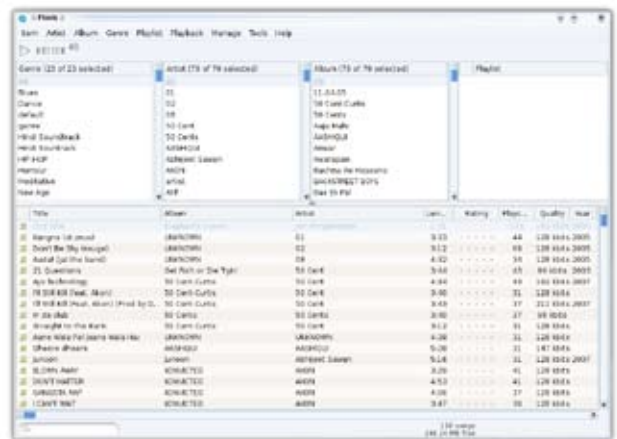


Figure 10: Floola iPod manager

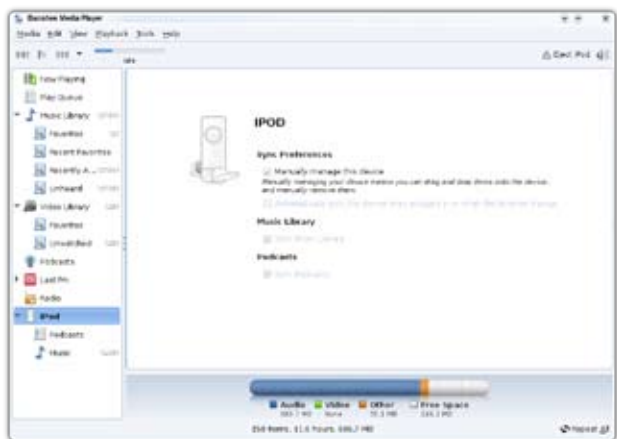



Figure 11: Banshee displaying iPod

dropping from Amarok or Banshee. iPod users can add media files from the suggested media player or they can use iPod managers. Linux has some great iPod managers like *gtkpod* (Figure 9) and *Floola* (Figure 10). All the iPod managers and media players will transfer album art along with media files.

That's it for now! Hope this article helps you to manage your music much more efficiently, allowing you to get the most of every new function that comes with the new breed of media players. **END** 

Resources

- Banshee Media Player: <http://banshee-project.org/>
- Amarok Media Player: <http://amarok.kde.org/>
- Easy Tag ID Editor: <http://easytag.sourceforge.net/>
- Picard Tagger: <http://musicbrainz.org/doc/PicardTagger>
- Banshee PPA: <https://edge.launchpad.net/~banshee-team/+archive>
- GTK POD: <http://www.gtkpod.org/about.html>
- Floola Ipod Manager: <http://www.floola.com/>

By: Shashwat Pant

The author is a FOSS enthusiast interested in Qt programming and technology. He is fond of reviewing the latest OSS tools and distros.



Slax 6

Slacks Off To You!

Slax Graphics mode (KDE)

Slax Always Fresh

Slax Copy To RAM

Slax Graphics UESA mode

Slax Text mode

Slax as PXE server

Run Mntest utility

More about currently selected:

Slax the best way we can.
 try to autoconfigure graphics
 and use the maximum
 allowed resolution

Automatic boot in 7 seconds...

First, there was Slackware. And then there was Slax. As the similarity between the names suggests, Slax is actually a size-optimised (well, from 1.9 GB worth of installation files to a 190.1MB LiveCD) version of Slackware that's meant for use as a Live CD and LiveUSB.

We have seen smaller Live distros than Slax (Slitaz, for instance), but Slax is by far the most famous and proven Live distro. Slax, though not at all officially related to Slackware, rigorously follows the Slackware release cycle. The subject of our review, version 6.0.9, was released on the same day as Slackware 12.2 (December 10, 2008). Built from Slackware 12.2, it has the same rock solid stability and simplicity, while maintaining the ease of use not found in its upstream versions.

Slax 5 to 6

Those who are still using Slax version 5 will find a lot of differences between Slax 5 and 6. These are:

1. Most importantly, version 6 is available in only one edition. Slax

5 had Standard, KillBill (Microsoft Killer), Foro, Popcorn and Server editions. Slax 6 has none of the KillBill or Server features; they need additional modules to be downloaded.

2. This brings us to modules. Previously, the Slax base was not a module, while anything running on top of it, was (such as X11). Now, everything is a module, even the base. Slax 6 comes with 6 modules—core, xorg, kde, kdeapps, devel and koffice. You can remove or replace these base components if you wish.
3. X starts automatically! There's no need to run *xconf* and *startx*. The downside of this is that you run in X as the root.
4. The Slax bootsplash flower has been done away with—there's only boring text now.

Go Slax, go!

I downloaded Slax 6.0.9 and made it boot up on a VMWare Workstation VM. It, as always, didn't fail to surprise me with its sub-10 second boot to KDE. But the boot threw up another surprise—a new option to start Slax as a PXE server. It turned out that this option makes Slax boot up as usual, but starts a TFTP server in the background with the Slax CD as the root. This means that if you were to start up a computer and boot from a network, Slax would boot as if the CD was inserted locally in the client's PC. Neat!

I still detest what they did to the bootsplash. Slax 5 used to have a Slax flower above the scrolling boot log, but it has been done away with in version 6. I wonder why.

Slax's 'Always Fresh' feature is another neat trick. Suppose you had been saving a persistent home all along, but suddenly wanted to boot a pristine distribution, the 'Always Fresh' option will ignore the persistence file.

You don't need to download the tarball and ISO separately for the USB and CD; each will work for the other. If you downloaded the ISO, burn it to a CD or DVD and it will work fine. If you want a USB version, copy everything inside the ISO to the root of the USB drive and run `/media/disk1/boot/bootinst.sh` (assuming `/media/disk1/` is the directory where your USB stick is mounted) to make the USB stick bootable. Conversely, for the tarball, the USB installation procedure is the same, but if you need an ISO, run the script `slax/make_iso.sh` to create an ISO and burn it to a CD or DVD.

The GUI—Slacks and boots

Less than 10 seconds after I hit *ENTER* on the SYSLINUX boot menu, I was greeted by the KDE 3.5.10 splash screen and the Boots wallpaper. Nothing has really changed here; it's all the same as in the earlier version 6 desktops. One word of advice: the root password is 'toor', if you ever need it.

X defaulted to 800 x 600 in VMWare. It doesn't include ATI or NVIDIA drivers, and due to a lack of blanks, I was unable to test out native on my NVidia 7100 iGPU. If it doesn't work, you should be dropped to a console. Type in *guisafe* and X should start up in VESA mode.

Table 1 sums up all the software included with this version. Slackware 12.2 came as a disappointment to me—no KDE4, no Python 2.6 and no GCC 4.3, while Slax seems to have taken all the software that Slackware had. The latest KDE 3.5 stable release is included, and all the apps are Qt-based (with no space for GTK+).

The absence of Mozilla Firefox really infuriated me; Konqueror 3 is not at all a replacement for Firefox, but it still works. When it comes to the Web, I was unable to configure my ADSL (broadband) service in a bridge mode whatever I did. Looks like I have to keep

List of applications installed by default	
Category	Software included
Games	KBattleship, KBounce, Patience
Graphics	KuickShow (image viewer), KolourPaint (Paint-like software), KSnapshot (screen-shot capture), KKolourChooser (colour chooser)
Internet	Konqueror, KMail, Kopete, Akregator, Krdc (remote desktop), Krfb (VNC), KNetAttach (network folder browser), KPPP, network-conf, KWiFiManager
Multimedia	JuK, KPlayer, KsCD, K3B, KAudioCreator, KMix
Office	KWord, KSpread, KPresenter, Kontact, KPDF
Utilities	KJots, KWrite, KNotes, Klipper, KCalc, Ark
System	Slax Module Manager, KInfoCenter, KSysGuard
Command Line	GCC 4.2.4, BusyBox 1.11.1

Table 1



Figure 1: Default desktop

my modem in PPPoE mode. RP-PPPoE is included, but it didn't work. Two more things to whine about—no torrent client and no firewall.

As with Slax, MP3s play out of the box. WMAs don't, but that's the same as with JuK. The media player KPlayer uses MPlayer as its backend and can play anything that is thrown at it.

Under the hood

Under the hood, Slax runs on Linux 2.6.27.8, which is not the stock Slackware kernel. Since the kernel had to be re-built (patched to support LZMA and AUFS), I guess Tomáš Matějčík (the Slax developer) went with the latest kernel. All the command-line utilities are provided by BusyBox 1.11.1 (what a version!). SquashFS has been upgraded to version 3.4 and fixes a bug with the earlier UNSQUASHFS version. Slax 6 uses

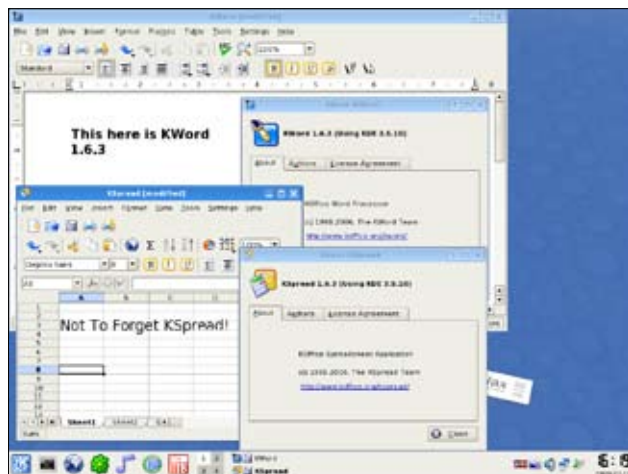


Figure 2: KOffice applications



Figure 3: K3b CD burner and KPlayer media player

AuFS (Another Union File System) to stack each of the SquashFS modules. It boots using the `initrd`, which has the AuFS drivers, then inserts each module from the base directory. When it's done, all modules from the modules directory are inserted. Finally, KDE is started!

What else?

Customised versions of Slax don't have to be built manually. You can get a custom build from the website. Go to www.slax.com/build.php. Add or remove any modules you may want. As of now, you can choose between 1,284 modules. When you are done, use the links to download a TAR or ISO. There are some limitations, such as non-resumable downloads, lack of personalisation options, etc, but for now, a rudimentary set-up like this works okay.

Derivatives

Slax is a very modular and layman-customisable distro. Thus it is no surprise that derivatives start springing up all over the place. Linvo [linvo.linux-bg.org] deserves a special mention. It is a GNOME-based Slax-derived distro that is by all means complete. It includes Wine, Firefox and Code::Blocks for development (they just hit my soft corner with Code::Blocks). Linvo can be used as a minimalist home

distro that includes one app for every task you need to perform.

The bottom line

Slax 6.0.9 is as always a heck of a distribution. It's by no means a distro you would use for recovery purposes; for that, SystemRescueCd is the one and only choice -- but as a minimalist distribution, it is very complete. One possible use is that you could create a Persistent LiveUSB with Slax and carry it around with all your files. This way you have an entire computer in your pocket, which just needs a... well, computer, to run. For basic recovery purposes (such as an OS crash or back-up restore) Slax can work, and you can use it as the main OS in lower spec'd or prehistoric computers (it just needs 128 MB of RAM to work). It's very uncomplicated, and can be used by beginners or teachers to teach Linux (from an end-user viewpoint). Overall, it is a worthy upgrade to any Slax fan, Slackware fan, Linux lover, Linux hater... well, just about everybody.

Postscript: Linux Live Scripts

Slax is built from a stock Slackware installation and a customised kernel. And the good news is, the entire build system is available for you to use and is distro-independent. So if you want to build Ubuntu-Slax, Debian-Slax or

Fedora-Slax, or just about any Slax, start by installing LZMA, AuFS and SquashFS. Then get the kernel source and patch it to support all of these. Finally, go to www.linux-live.org and download the live scripts. Execute them to get your own Slax ISO. All instructions are available in detail at the Linux-live website. Remember, by this method, anything in your system's filesystem is included in the Live distro.



Slax 6.0.9



Pros:

Minimalist yet complete, fast boot-up time, easy to use, lots of boot-up features, small size. Oh yes, it works faster than a Saturn V rocket.

Cons:

No updated software, no bootplash... well, actually nothing serious to complain about.

Platform: x86 (No separate version for x86_64)

Price: Free (as in beer)

Website: www.slax.org

By: Boudhayan Gupta

The author is a 14-year-old student studying in Class 8. He is a logician (as opposed to a magician), a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++, Python and Assembly (NASM Syntax).



SICSR
Symbiosis Institute of Computer
Studies and Research



13th and 14th
February, 2009

GNU Unify⁰⁹
A Forum to Unite Open Minds

The GNU Unified Experience!

Here's a report on an event that touched on almost all aspects of open source—from installations to kernel programming and scripting, from OpenOffice.org to Bash and Perl, from Ext4 file system to SCSI, and from scientific computing to network security.

*G*NUUnify '09 at SICSR (Symbiosis Institute of Computer Studies and Research) on February 13 and 14 was a 'dear diary' moment for open source enthusiasts in and around Pune. As GNUUnify entered its seventh year, PLUG (Pune Linux Users Group) and Mozilla joined hands to make it a big success.

The two-day action-packed event had multiple activities running in parallel. The program was carefully designed so that there was a solid takeaway for different sections of audience—from engineering students to desktop users and from administrators to elders trying to switch to open source. Let me get down to sharing my '*GNUUnified experience*'.



▲ The full house



▲ The Mozilla camp in action



▲ Saifi Khan talks about open source firewalls



▲ Dextor in a workshop



▲ InstallFeast: helping elders in the installation process



▲ Participants grab a quick bite

Day 1: Programming, storage, networking, et al.

Day 1, which was February 13, started with workshops and four parallel tech tracks. Manjusha Joshi led a workshop on the TeX/LaTeX documentation tool, followed by Rajesh Sola's hands-on tricks and tips on OpenOffice.org. Ebenezer and Samar also held a workshop in another classroom on *Tweaking Ubuntu*, wherein they showed us how to configure Ubuntu, connect to the repository, *apt-get* the required packages, create an ISO image, etc. They also demonstrated VirtualBox and how to run Ubuntu on Ubuntu.

The scheduled talks from Day 1 could be broadly classified into three streams:

- The programmers' track
- Storage and networking
- The PHP marathon

To start with, I attended a talk on open source firewalls by Saifi Khan. It mainly covered how a firewall is not just an application software running on a single computer but has to deal with the complete network infrastructure. He explained the support available in FreeBSD, NetBSD and Linux for writing firewalls, and followed that with a talk on *iptables* usage with practical demonstrations.

Next, I listened to a short but interesting talk about the Generic Netlink Socket framework by Alok Barsode. He explained this powerful IPC technique in detail and this was followed by a practical demonstration of a driver written using the Generic Netlink framework.

The next talk I attended was equally exciting. Ajay Kumar, winner of the Google Summer of Code 2008, spoke on the humanitarian FOSS project dubbed 'Sahana', explaining the technology involved and the challenges it faced.

I then attended a session on an open source library management system called 'Koha'. Krishnan Mani started with how he himself ended up getting Koha up and running for his community library. He also talked about the deployment of Koha in India

and abroad, and various techniques to migrate to Koha from existing 'spreadsheets'.

Being a kernel enthusiast myself, I next headed over towards the kernel discussion panel organised by *GeeP-Geeks Of Pune*. Linux kernel gurus like Amit Shah, Amit Kale, Anand Mitra and Kedar Sovani chaired the panel and addressed questions like how different kernel programming is, how to start off with it, what the challenges are in debugging, how and where to get support and documentation from, etc.

On the same day, I also got a chance to peep into an 'Install Fest' activity for the Fedora core Linux distribution. I found students as well as elders trying out Fedora installations and coming out with a degree of confidence.

Day 2: Mozilla, Fedora and more programming

Workshops on the second day consisted of a hands-on session led by Abhishek Nagar. Aligned with the PHP marathon, as the name '*Fast track websites—from local to remote*' suggests, the workshop was based on how to build a website using Drupal. Running parallel, Vinay Pawar (a.k.a Zoid) led another hands-on session on 'Blender', making people think in a third dimension.

The scheduled talks from Day 2 could again be classified under:

- A programmer's track
- The Mozilla Camp and networking
- The PHP marathon (continued) and a Fedora activity day

Being a programmer, I chose a talk by B.C. Sekar on *Doing Linux Projects*. The talk focused on the power and benefits FOSS offers as a project development platform. He also introduced the audience to the basic licences, such as GPL and LGPL and other things a developer needs to know before contributing to an open source project. He shared his experiences in enabling his customers to achieve an improved time-to-market, using off-the-shelf FOSS tools as against the lengthy

process of developing proprietary tools.


Next, I got to look in on the Mozilla Camp. Seth Bindernagel and Arun Ranganathan co-hosted the camp. It started from the history of Mozilla and how it was born from Netscape, and progressed to quite an interesting and interactive discussion focusing on the recent release of Bepin, followed by a demonstration on Bepin.

I then got to listen to Alolita Sharma, who focused on one distinguishing characteristic of open source—of users becoming the contributors and the resulting decentralisation of ownership, synergies between the people of different cultures, the give and take of polite feedback rather than blame games, etc. The talk was also backed up with case studies of three important open source projects—WordPress, Ubuntu and Mozilla Firefox.

The talk titled *FREEconomics: The economics of free/open source* by Navin Kabra was also insightful. As the name suggests, it was all about the business model of FOSS-based companies. He highlighted the characteristic economics behind 'free and open source software products'. He concluded by putting forward an interesting and important idea about the 'Attention Economy/ Reputation Economy'.

The last talk I sat through was on '*My experience with Linux as a customer*' by Atul Tulshibagwale. The speaker focused on the gaps in the current Linux distributions, which make the Indian audience hesitate to use it as their desktop OS. He emphasised the important applications that still need to be ported, and how the small but smart changes in default configurations of applications like OpenOffice.org would make the transition comfortable for the Indian audience.

Finally, before leaving for the day, I sneaked a peek into the Marathi Localisation activity room. What was in progress was a session on '*FUEL-Frequently Used Entries for Localisation*', by G Karunakar and Sandeep Shedmake. I noticed a number of Marathi wits translating and verifying the 'correct and closest' translations for the different terminology in OpenOffice.org, making it easy to understand for local people. This was a two-day ongoing activity and covered more than 250 terms (as per the last word count that I heard about). I contributed a few and left for the day.

All in all, I'd say the organisers and volunteers worked really hard to make sure the event was yet another success story in the history of GNUnify. By the way, have you visited gnunify.in yet? 

By: Nilesh Govande

The author is a Linux enthusiast and can be contacted at nileshgovande@yahoo.com. His areas of interest include Linux systems software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.

PLEXUS

Hardware & Networking

Courses Offered

RHCT

RHCE



- Excellent Infrastructure.
- Certified Faculties.
- New batch starts on every Monday.
- Register for RHCE exam and avail **5 DAYS FREE** intensive examination coaching .

Plexus Software Security Systems Pvt.Ltd.

#31/14, Burkit Road, T-Nagar
Chennai-600 017.

Ph:044-2433 7355 Mob: 99401 66510

Email:training@plexus.co.in

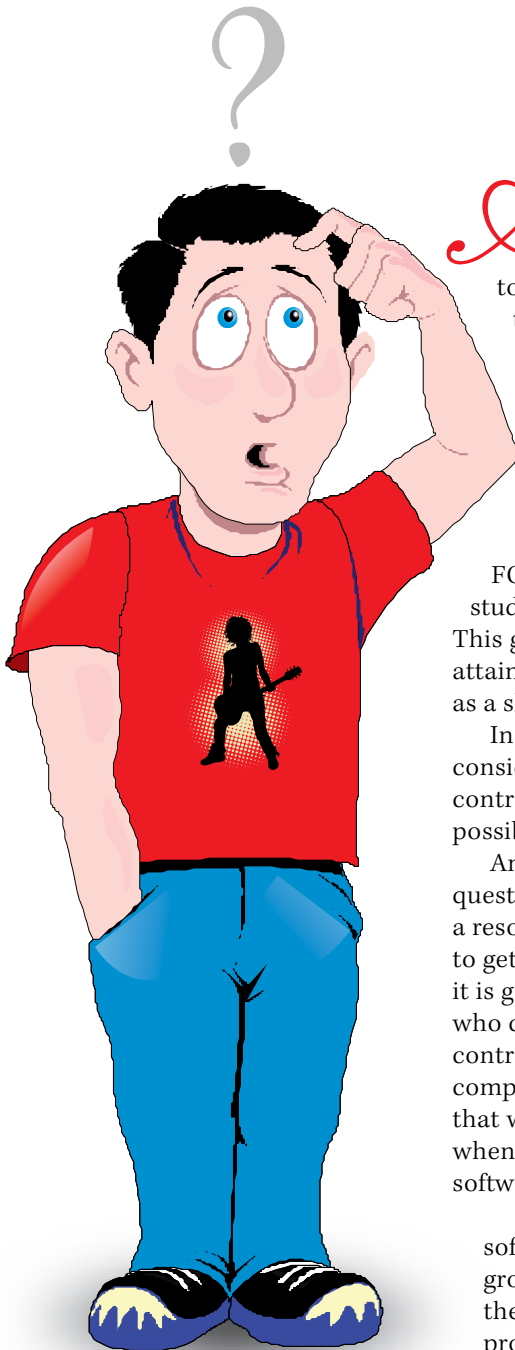
URL: www.plexus.co.in

Other Technical Partners



Will FOSS Get Me A Job?

FOSS allows anyone to acquire the skills that lead to becoming a better developer and an improved person.



Any introductory talk on Free and Open Source Software (FOSS) addressed to students will throw up the typical question: “Will FOSS get me a job?” This is generally a follow-up question to “Why should I do this FOSS thing?” A lot of blogs and articles that I read state that in the current economic downturn, FOSS ought to be something students should be looking at. This goes to prove that FOSS has attained mainstream acceptance as a skill worth acquiring.

In short, students should consider participating in and contributing to FOSS as early as possible.

And, as a response to the first question, the answer is generally a resounding NO. FOSS isn’t going to get any student a job. However, it is going to equip anyone who chooses to participate or contribute with the required skills, competence, and the recognition that will surely come in handy when building a career around software development.

Training a recently recruited software developer from the ground-up in the basics of the software development process is an expensive and

labour-intensive affair. Yet, lots of companies do so because of the lack of such skills in their freshers. This is one aspect that can be taken care of by gradually acquiring skills in the world of FOSS. The academic curricula that students go through bring them up to speed with the rigour and discipline imposed by the theories. FOSS allows them to immediately implement their knowledge and learn from a collaborative experience.

Let’s take the example of students who feel motivated enough to begin by participating in and thereon moving to contributing to FOSS. What skills do they obtain? Plenty.

Since FOSS development is mainly driven over the Internet, the very first skills that get polished are communication skills and the ability to use communication tools like e-mail and IRC (and IMs). Virtual communication puts the responsibility on the sender of the message to be clear, concise and precise. All these are very good qualities to be learnt. Additionally, appreciation of the cultural nuances of interaction, the social norms, etc, make a new contributor a much more well rounded personality in addition to enhancing developer skills.

Moving on, any FOSS project


would have its version control system, and submissions of code or content to the version control system undergo the age-old process of peer review. How are these two important? It helps a student get familiar with the theory and practices of version control, the need to write code/content/patches according to established guidelines, and build upon the communication skills learnt to appreciate the feedback from a peer review group.

So, from just interested participants, willing students are well on the road to becoming well-rounded developers with various skills that make them invaluable when the recruiting season comes around. But wait, there's more: FOSS development processes ensure that contributions of code/content are always out in the open and available for perusal/analysis. What this means is a portfolio of development work. How does that help? Well, if there is an existing body of peer-reviewed code/content on a publicly-available version control system, it helps a recruiter do a technical assessment of the candidate. This does not really mean that a company would waive standard procedures of technical tests, but it would perhaps be of an added advantage when put in the perspective of peers. And, for companies already doing FOSS, such a code/content portfolio is of immense use. It allows them to form a judgement around the competencies of the candidate and even check out with the project module leads about how good the contributor is.

The curricula teaches the students about the Software Development Life Cycle (SDLC), its various stages, and the different checks and documentation that make up the Body of Knowledge. Participating in an upstream FOSS project provides an excellent exposure to these intricacies.

Following up a project roadmap and working on tasks related to modules/components, while keeping in mind the project release cycles, allows the contributor to become proficient in the real-life aspects of the SDLC concepts. Additionally, with time and an increase in the quantum of contributions, the new contributor would soon be confident enough to help out and mentor others in the project. Thus completing the circle, while learning how to work with dispersed teams, communicate virtually and work to timelines.

These are qualities that companies spend an inordinate amount of time inculcating in their new recruits; while participating in a FOSS model of software development, anyone can learn it as on-the-job training. This can be done in addition to the activities of an academic life. Lessons from books are somewhat easily tested and applied in real-life projects. Thus, students should take the time to look at any interesting project and make an effort to participate in it.

So, why did I say that FOSS will not get students a job? It should be fairly obvious by now. Brushing your teeth regularly does not automatically make you a film star, does it? But good dental hygiene along with disciplined practice should equip you with a pleasing personality that may (or, may not) lead to stardom. In somewhat a similar way, FOSS allows anyone to acquire skills and personality traits that lead towards becoming a better developer and an improved person—which is a long way down the road towards building a good career. **END** 

By: Sankarshan Mukhopadhyay

The author has been using Fedora since the days of Fedora Core 1. He can be reached at morpheus at fedoraproject dot org, or as sankarshan at jabber dot com on IM.

p00r JOB MARKET ?
Campus
Placement ?

Future is OPEN

Now, time to
Upgrade your Skills. Register...



RHCE

Free Training for RHCE Exam
RHCE Exam - 17 & 30 Mar 2009



Authorized
Training Partner

First Time in S India
Only at

Zend Certification

ZCE - International Certification for PHP
from Zend Technologies - USA

**International Certification
for PHP (LAMP)**

Zend Certified Engineer (ZCE)

PHP I - Foundation

PHP II - Higher Structure

PHP 5 Certification

Save
Upto
50%
Now

Offer Valid upto 20 Mar 2009



Future is open

Lynus Academy Pvt Ltd

Chennai - 044-25021630,
42171278, 9952966527

Pondicherry - 0413-2276233,
4202121, 9952966528

contactus@lynusacademy.com
www.lynusacademy.com

• Conditions Apply

The Open Movement and the Implications of its Opportunities in Education

In this article, we look at what role the various 'open movements' can play in academia, and how academia should nurture them for a win-win relationship.

Over the last few years, the concept of openness has been spreading its wings far and wide in many guises. Much of it started with the popularity of the FOSS (free and open source software) movement. Though dating back to Richard M. Stallman's days at MIT, it gathered popularity only on the arrival of the Linux kernel from Linus Torvalds in the early 90s. Since then, the movement has never looked back.

Almost every large corporation is involved in the movement in one way or another, with IBM, Sun, Intel, HP, etc, leading the way. The movement has got into the legislative bodies of many countries, creating pressure at the level of policies, guidelines, etc, to support and adopt FOSS wherever possible. Many countries have taken explicit steps to nurture a FOSS ecosystem, by training programs, certifications, resource centres, and so on.

Today, FOSS is a familiar term across

the world, spanning academia, industry, government, and SMEs. The FOSS activities have a number of dimensions, stretching well beyond the availability of software with source code. It has led to sibling movements in content, standards, hardware, systems, etc. Each of these has attained a fair degree of maturity today.

Academia has always related to FOSS's way of thought quite easily, thanks to the same underlying philosophy. However, there have been pressures from the proprietary software world, often derailing the curriculum in many cases, and disrupting the balance between conceptual foundations and commercial aspects. An example is the issue of a vendor-neutral syllabus in India, which has been in the air for a long time.

Here, we look at the role the various 'open movements'—a term to denote the set of movements consisting of free and open source software, open standards, open content, open

hardware, etc—can play in academia and how academia should use this effectively to nurture these movements in return, for a win-win relationship.

We will first discuss the basic driving force of a common mindset, and the associated implications. Then we look at specific aspects such as content, standards, software/hardware, etc. Though licensing is a major issue in this regard, we will ignore the licence discussion in this article, it being more appropriate for someone with a legal background to talk about this.

Philosophical match

Academia is about sharing knowledge—to enrich the giver and the receiver. One builds on the knowledge created by others, and shares the enriched knowledge with others to let it grow further. In a broad sense, software and content can be seen as embodiments of knowledge, differing, perhaps, in the way the knowledge is captured and represented. And hence, the notion that software needs to be available with the full source code is something natural to the academic community.

It was the commercial companies that introduced the concept of hiding the source behind a compilation process, to ‘protect’ what they call as their ‘intellectual property’, in turn to ensure that copies (and even more so, any modifications) are not done by the customers. The matching mindsets have many implications for both academia and the open movement.

The popularity of open source is the highest in academia. The pricing issue is certainly a major factor here, since academia often has the severest budget constraints in acquiring resources such as software and equipment. The ability for academia to understand the philosophy behind the open movement, in the sense of sharing knowledge freely, also plays a major role. Their contribution in pushing the quality and quantity of open source higher has been significant.

Researchers, even earlier, used to contribute the software developed for their research work, often in cutting areas of development, to public use. The Moodle learning management system, Latex document processing system, etc, are good examples of high-quality systems coming out this way.

Another implication of this link is the growth of open standards. FOSS is strongly based on the community metaphor of the *bazaar* development model described by Eric S. Raymond (ESR), which brings together a number of people from around the world to work on a common system. The roles for each are open, and how much they contribute to the final system is also open. There are only internal deadlines set by oneself.

This necessarily demands explicit efforts to reduce the learning curve for others, and transparency in shared data structures. Open standards, where the complete description is freely accessible to everyone, and where the standards evolve from collective contributions, becomes a natural choice. Not surprisingly, most of the open source programs use open standards wherever available. Here, formats and conventions invented for one system are reused for another, if relevant.

This brings us to the idea of building on what is existing—another typically academic mindset. As Newton remarked, “If I have seen farther, it is by standing on the shoulders of giants.” Research literature that builds on earlier results and acknowledges them by citation, and new software programs that reuse and extend existing software programs are clearly echoing the same mindset.

Starting from scratch every time, does not take you very far, in trying out new ideas in an already rich landscape. Starting from something that provides a close approximation to what you are looking for, and extending it, as appropriate, is more productive. As ESR has remarked: “A good programmer knows what code to write; a great programmer knows what to rewrite.” This kind of reuse and extension necessitates the openness associated with the open movements. In fact, all these characteristics are fundamental to the growth and sustenance of open movements. And the outcome of these movements, in turn, contributes substantially to the growth and effectiveness of academia.

Open content

The last few years have seen tremendous growth in open content, where the content is declared free for use, just as was done for FOSS. Comparable to FOSS licences like GPL, LGPL, etc, a group of licences also evolved to provide legal sanctity to this move. These are known as Creative Commons and offer a family of licences embodying the core idea of openness, and providing options for permission to modify, retaining attribution, and so on.

Perhaps, the most classic example of open content today is Wikipedia, which takes the openness to the extreme, allowing anyone with an account (that can be freely obtained) to modify any of the pages. However, the quality of the content on Wikipedia is generally very good, and some formal studies have also shown this empirically. Though there are topics involving strong controversies, where one often sees a series of continuous modifications by the opposing sides to support their stand.

For most academic content, the Wikipedia offers excellent reference/learning material with additional links, images, and so on. Lacking even a core group to filter modifications as is done with open source software, the high-quality of content indeed shows the feasibility of the approach.

The EU-funded SELF project exploits such resources, to even form course material for university courses, dynamically. Other examples of open content are the ‘million books’ of Project Gutenberg, the audio book collection at Librivox, the number of video repositories of Google video, YouTube, and so on. The movement got its momentum from the MIT initiative of open course ware, which has, in turn, led to a wider initiative of the open knowledge initiative (OKI) involving a number of partner institutions to share such resources.

It may be noted that these different set-ups follow different norms as far as their policy of use and modifications are concerned. A lot of open source software documentation and learning material are also available in such free content. These

Some FOSS educational software	
Application	Purpose
Euler	Complex numbers and matrices
Kstars	Astronomy with over 1.3 lakh stars, all planets, etc.
Chemlab	Chemistry lab
Sage	Algebra, geometry, etc
Units	Unit conversion
Earth3D	Real-time 3D view of earth
Kalzium	Periodic table and properties of elements
Atomix	Puzzle game for physics
Kig	High precision geometric constructions
Xaos	Fractal geometry

Table 1

include, machine learning with *Weka* (full book available online), O'Reilly publications, *NL Toolkit* (full book on this available, along with the tool kit), the Linux Documentation Project, and so on.

FOSS for education

It is in the area of software resources that the open movement has contributed most to the education sector. The software relevance to education is from three different angles, and these are described briefly in the following sub-sections.

FOSS learning resources

E-learning is another buzzword that is popular among all academic communities, though its meaning and adoption varies widely from group to group. One major concern in e-learning is the quality of content. Traditional content has been largely text and static pictures/images, limited by the medium of textbooks. Much of e-learning content is still restricted to these two.

Animations, simulations and interactive problem solving environments (IPSE) can significantly enhance the teaching learning process. It provides an opportunity to use the multiple senses in absorbing a concept, and also to try out the concept in perhaps restricted environments, through simulation and IPSEs.

These are generally hard to develop, as they involve a significant amount of software development for each topic. The system needs to have a fairly sophisticated model of the content relating to the domain, and be able to recognise and react to the events with respect to the domain. For example, a program illustrating the concept of projectile motion needs to be able to compute the path of the projectile based on relevant parameters—the initial velocity, the angle of throw and gravity. As these parameters are varied by the student, the system needs to revise the computation accordingly. These tools make e-learning much richer than what is possible in a traditional environment, and ought to be part of e-learning.

One reason for the ineffectiveness of e-learning in academic settings is the lack of such quality content, which would deepen the learning and encourage students

to use these. A lot of high-quality programs of this type are available in open source over the Web. Unfortunately, there are no reliable comprehensive repositories for these kinds of programs, as they are scattered efforts from people around the globe.

The UNESCO portal for FOSS, and repositories like the Edubuntu package list, provide some starting places. The OSCAR project of IIT-B also makes an attempt to collect animation programs. IPSEs are not included here, since these are often fairly large programs. Table 1 provides a (very small) sample of resources one can find on the Web.

One major challenge in using these resources for education is the need to link them explicitly into the curriculum. Except the highly-motivated students, most would be lost when exposed to these tools as a collection for them to explore on their own. For the purposeful use of these systems, activities assignments, experiments, etc, need to be formulated, using these tools.

FOSS for basic utilities

This is what's most obvious to people. Today, open source solutions of good quality are available for you to set up a basic computer system, without investing in any proprietary software. All the software components, including the basic operating system, office suite (for documents, presentation, drawing, equations, etc), browser, media players/editors, drawing utilities, network management, and so on, are available in open source today. In most of these cases, one also has a decent number of alternatives to choose from. Table 2 includes some of these tools. Installation and management of these are just as simple or complex as alternative proprietary solutions. FOSS-based desktops are seen to be generally less vulnerable to security problems such as virus infections; this is a major headache for systems administrators in educational establishments, in general.

Full systems customised for the educational sector are also available from some of the popular distributions. Examples are Edubuntu from Ubuntu, Eduknoppix from Knoppix and the proposed EduBoss from BOSS. These include the basic operating system and associated utilities, select tools for educational use (like an equation editor), and some learner resources for specific subjects. This removes the effort of having to pick the relevant packages from various repositories and integrating them individually.

FOSS for learning management

Under this category, I include software that is specifically for the educational environments. There are software solutions for school/college administration, for faculty to run and coordinate the various activities in a course, for faculty to create and manage content for a course, for students to track the progress and collaborate with other students, and so on. Accordingly, here too, the scope of software is vast, and FOSS doesn't disappoint us. Table 2 has a small list of some of these tools.

One can see a wide range of systems, from learning

management and school administration, to library management. Mostly, the systems listed are quite popular, with a good development and user community, and the software are quite stable and feature rich. Systems like Moodle have a large national and international user base. Many of these are also available in multiple Indian and international languages.

Future outlook

We can see that there is a strong synergy between the open movements and academic education. There is a lot that open movements are bringing and *can* further bring to academic programmes. We need to encourage our academia to benefit from this and also contribute back to help the growth of the movements, in return. Our own projects, students' projects, as well as PhD/MS projects can benefit tremendously from the existing resources, and can be used to drive new developments and significant enhancements. This needs to happen on a larger scale.

At the same time, there are new challenges coming up on the education side. Content sharing across institutions has happened relatively less often, so far. But with the growth of e-content and the increasing presence of institutions on the Web, this is bound to occur more frequently. Inter-institutional collaboration in sharing not only content, but full courses, subjects and faculty is certainly possible soon. This will necessitate a lot of changes in the software requirements, and offers good opportunities for us to contribute and also adapt existing software to meet these new requirements. New demands will also be made on interoperability as records and resources move across institutional boundaries. Initiatives such as OKI are a step in that direction. Work on distributed learning management systems also looks at similar concerns.

The sharing, in turn, also brings into focus the growing concern of plagiarism. With a vast collection of resources freely available, the chances of plagiarism and the difficulty in detecting it, is increasing. Scigen is a relevant case study, which produces 'scientific' research papers using some natural language processing techniques. Since the language appears of good quality, rich with a high degree of relevant jargon and following the style and conventions of a research paper, it appears genuine, and outputs from this program have been accepted in some international conferences. Tracking copied (with and without distortion) submissions for assignments to research papers is a major challenge in the academic environment.

There are also changes at the libraries. Good quality open source solutions are available to handle the functionalities of today's libraries. Even digital libraries are well supported by FOSS solutions such as Dspace. However, with the growth of e-content rich with simulations, and e-learning growing in popularity, the nature of resources that the library needs to deal with is changing. Dealing with IPSEs offers different challenges compared to conventional or even digital books.

Various FOSS applications by category	
Application category	FOSS applications available
Web browser	Firefox, Iceweasel, Konqueror, Epiphany
Document creation	OpenOffice.org, KOffice, Latex
Audio record/edit	Audacity, Ardour
Web page creation	Nvu, Bluefish, Quanta Plus
Content management	Drupal, Joomla, Plone/Zope
Learning management	Moodle, Sakai, Atutor
Question banking, testing	exe2learn, Moodle
School administration	schooltool
Visual programming	scratch
Diagram editing	Dia
Scanner	Xsane
3D animation	Blender
Image editing	The GIMP, Krita
Page layout program	Scribus
Plotting	Kmplot
Creating and running tests	Keduka
Video conferencing	Dimdim, Vmukti, Ekiga, open-meetings
Library management	Koha, Dspace

Table 2

Already, issues of subscription management and sharing of e-resources is a major concern.

We also need to look at removing the linguistic and physical barriers preventing people from using technology. Software localisation and accessibility are two fields related to these two aspects, which need to see a lot more activity, for countries like India.

In summary, FOSS has enriched the education field in many ways. But the world is moving fast in the education sector as well as other sectors, and new demands and opportunities are constantly emerging on the horizon. FOSS needs to be sustained and nurtured through a sustained cycle of human resources and efforts, to help it continue what it has been able to do so far. **END** 

References

- Good introductory material on open source, open standards, FOSS and education: www.iosn.net
- UNESCO portal on FOSS offering a list of open source resources as content and general software of use in education: www.unesco-ci.org/cgi-bin/portals/foss/page.cgi?d=1&g=10
- Eric Raymond. The Cathedral & the Bazaar. O'Reilly Media Inc, 2001. [Information on FOSS philosophy, development insights, etc.]

Dr M Sasikumar

The author has been with CDAC Mumbai (formerly NCST) for over 20 years and currently heads its artificial intelligence, e-learning and open source divisions. He can be reached at sasi@cdacmumbai.in or at the.little.sasi@gmail.com



How applicable is the open source software development model for e-governance initiatives?

Discourses on open source and free software generally approach the subject as a philosophy. There are a significant number of articles and papers that discuss free and open source software from the perspectives and interpretation offered by political economy and social commons. A suitably detailed perspective of open source was provided in *The Cathedral and the Bazaar* by Eric S Raymond. Since then, there have been sporadic attempts to reconcile the principles of open source with established practices in software development.

However, in recent times, there has been a renewed trend to strip away the theoretical bulwarks of the open source and Free Software movements and focus on the deliverables. Reinterpreting a few of those concepts, we would like to discuss the relevance of open source as a 'software development model' and its applicability to e-governance initiatives.

Open source is, primarily, a software development methodology. As is the norm with any other methodology or practice, open source has its unique ethos, rituals and codes. Hence, there is an established philosophy around it. Leaving the philosophy aside, the important aspects of open source as a software development method are:

- Collaborate to develop code
- Reuse innovation
- Release early, release often

We will use the above three aspects and derive the related ethos of open source and, in conclusion, will demonstrate the suitability of this model as applicable to ICT4D application development.

Open source mandates a level of intense collaboration. Such engagement between developers and users leads to an improvement in the quality of the code. Feature enhancements and the identification and resolution of defects occur seamlessly because the source code is available for

perusal, and is backed by a set of tools that allow the reporting and tracking of requests and bugs.

Since the first principles of software engineering are somewhat generic, an intense level of collaboration is reflected in the ability to reuse innovations at an intra-project and inter-project level.

If you collaborate to innovate results in short development sprints, it ensures that a prototype can be released early. And, since development happens within the boundaries of predefined milestones, releases happen often. This ensures that the end users have access to gradually improving versions of the application, with direct inputs into the development lifecycle and in feature enhancements.

So far so good! But what makes all this happen? Infrastructure and communication.

The primary infrastructure requirement that needs to be planned, prepared, configured and deployed, is a 'forge' or a collaboration platform. By providing the features of version control, forums, group mailing (or, mailing lists), issue tracker, etc, a collaboration platform provides an ideal way to initiate constant communication. Having a central repository of codebases (i.e., the source code along with means to tag and search) makes it easy for application of design patterns on code. It also facilitates peer review and thus leads to improved code strength.

A forge provides the ideal set of data points that make it possible for teams working on modules and projects to meet, brainstorm and discuss—both in person, as well as virtually. The world of open source makes it very normal to hold code review and stand up meetings virtually, using virtual whiteboard technologies like XMPP, Obby, etc. Coupling such communication infrastructure units with an established policy for software development makes for a win-win scenario while adopting the open source model.

The government of India has initiated a large number of citizen-centric IT initiatives with the objective of making available services to the common man. These initiatives are based on IT design patterns that share commonalities. Identification of such 'common' aspects and adopting an open source software development model using the appropriate infrastructure would lead to a set of innovation patterns.


The open source model can also be extended to projects beyond the citizen-centric initiatives and at the level of the electronic 'mart' projects being undertaken for agriculture. Since the model allows collaboration, it brings in an additional number of software developers towards reviewing code and fixing issues. Additionally, it allows for rapid prototyping, thus providing benefits of early release and testing of software. The model also lends itself well to issues of standard definition, font creation, auditing of Web services for security and standards compliance, among other things. The underlying theme is to not limit the open source model towards software development, but to extend, adapt and adopt it to as many aspects of the work

flow as is permitted, with the end objective of producing high-quality software and content.

The open source model for software development is not limited to putting in place technology infrastructure—hardware and software. It also requires that an adequate definition of the software development policy is in place. Such a policy would include methods to check and control, that provide the project administrators with a granular view of the project's progress. Additionally, this policy would encourage code review, code reuse and enhancements to existing codebases.

The reason for such a model to attain a measure of success would be because it does not mandate a complete overhaul of the existing process. The open source model for software development puts in place a system of project management that allows innovation to be transparent, prototyping to be rapid and collaboration to be constant. A resultant effect of approaching software development using aspects of open source is a higher degree of collaboration by leveraging the effect of 'crowd-sourcing'. In effect, doing code development in the 'incrementally improving' method allows for a greater degree of rigour in using test cases. Such a course of action would create improved software. The ability to prototype rapidly also provides the scope of customisation.

The constant availability of a 'forge' or repository of projects and associated codebases provides a unique opportunity to estimate the maturity of a software development unit. The added advantage of reuse allows a metadata component to be added to such estimations. It is somewhat logical that robust codebases and strong design patterns would be adopted across the organisation. However, such adoptions do not lead to a monoculture of coding patterns. Instead, it allows creativity to flourish within the scope of coding standards and project guidelines as defined by the policy formulating entities. The repository provides a replacement to unorganised institutional memory and, in a way, provides a historical timeline of projects as they evolved through their project management charter.

To sum up, adopting and adapting the open source model of software development for government-funded projects provides a win-win scenario for all stakeholders. The developers are attuned to the idea of rapid prototyping and release, which makes for an increased involvement from the consumers/customers of the software. And the resulting feedback loop increases the robustness, efficiency and innovation factor of the software projects. Collaborating to innovate provides the bedrock for a much shorter time-to-release in the software development lifecycle of a software project.  **END**

By: Sankarshan Mukhopadhyay

The author works at Red Hat and has been using Fedora since the days of Fedora Core 1. He can be reached at sankarshan@redhat.com or at sankarshan@jabber.org (instant messaging).

With simplicity and stability continuing to be top priorities,

Slackware 12.2

doesn't disappoint.

I have tried out many flavours of GNU/Linux over the years—starting from Red Hat 9 to Fedora 10, Ubuntu, Sabayon, openSUSE, etc. In my book, all of these have their pros and cons, although I keep distro hopping not because of their cons, but because I am fickle and get bored with using the same thing for a period of time.

When I heard that *LFY* was bundling the latest Slackware 12.2, it reminded me that this was one distro I was yet to try out; besides, I wasn't all that happy with Sabayon 4 that I was currently using on my home computer. Slackware, as you know, is the oldest surviving GNU/Linux, whose roots go way back to the early 90s, and a motto to keep things stable and simple. Version 12.2 was released on December 11, 2008, and I wanted to see if they still lived by that motto.

Now, when Slackware says that it's focus is on simplicity, don't take that to mean the typical 'click-next' stuff we generally associate the term with. In fact, technically speaking, 'click-next' makes things more and more complicated due to various top-tier UI abstractions that try to obscure the backend command-line. So, when they say Slackware likes to stick to simplicity, it means things are kept

in line with the upstream, without any distro-specific customisation. In a way, it tries to stay aligned to the original UNIX philosophy—why break things if it still works well.

This brings me to the installation procedure. You'll still find the same installer that's perhaps been in use for the last few years—maybe even for a decade or more. The installation and package configuration for Slackware is all text-based.

Before I talk about the installation, a new experience for me, here are the specs for my test system

- Pentium 4 with HT technology
- Intel 865 motherboard
- 1 GB RAM
- 80 GB hard drive

Ready for installation

After popping in the DVD, I was greeted with a non-graphical interface. A word of advice for people who have never tried anything that's not graphical: you need to pay attention to the things being displayed on the screen... keep reading!

Now, with the bootable DVD inside and everything on-screen in black and white, I started reading every page on-screen to make the proper choices. First, you will be asked to select a keyboard layout. Once

STABILITY

Out of the Box

done, you will be given a root shell prompt. From here, you will need to execute the correct commands and make a proper partition selection to install Slackware—remember, you generally don't have those *Back* buttons as in the case of the command line. This is the most complicated part—a single mistake here can format your earlier data. So, be very cautious while creating the partitions for installation.

fdisk or *cfdisk* are the two commands that can be used to partition the hard drive. I did it using *fdisk* because I am familiar with it—although *cfdisk* is more user-friendly, they say. Once partitions are created, you will have to type *setup* and hit *Enter* to begin the installation.

Next, it will ask you to format the partitions that you have created and you can also mount the other partitions in Slackware. The other non-Slackware partitions can be given permission like read/write/execute. Once you are done with all these, you can select CD-ROM as the source media for installing Slackware into your hard drive. You should then get an on-screen prompt asking for an auto (recommended) or manual device name (CD ROM) selection.

Slackware will also give you an option for selecting the packages, like *full*, *menu*, *expert*, *newbie*, *custom*, *tagpath* and *help*. I went with *full*, and the package installation started. The full installation takes around 5 GB of disk space. Anyway, it was time for me to go get myself a cup of tea.

After the package installation is over—the procedure on my system took around 50 minutes—an on-screen text menu appears asking whether you want to use a USB device to boot Slackware. I chose to skip this option. Now comes the most critical part of the installation, the boot loader. In Slackware, the bootloader is still LILO—I told you, they don't change things if they are not broken—so no Grub. You must be very careful while installing LILO. During its installation, you are given three choices: simple, expert or else you can skip and install LILO later. I opted for the 'simple' option and then I chose the standard resolution for my screen. You will also be asked about the location to install LILO (root, floppy, MBR). I chose MBR. After that, you're prompted for the mouse, network, font configurations, hardware clock settings, time zone selection and default window manager for X. After the entire set up is done, you get a prompt saying 'Setup complete'. Now exit the installation and reboot your system by entering *init 6* from the root prompt.

The system boots into a text mode. In order to view the graphical interface, you will have to log in as the root and type *startx* or *kdm* to start the graphical desktop. I hit *kdm*, as *startx* would log me into the GUI directly as the root. On the KDM screen I released I have to create a user account first.

So, it was time to go back to the root shell by hitting *Ctrl+Alt+F1*. First, execute the following to create a user account:

```
# useradd -m abhijit
```

Abhijit happens to be *my* name. Remember to replace that with yours. Now, set the user account password with



Figure 1: The default desktop

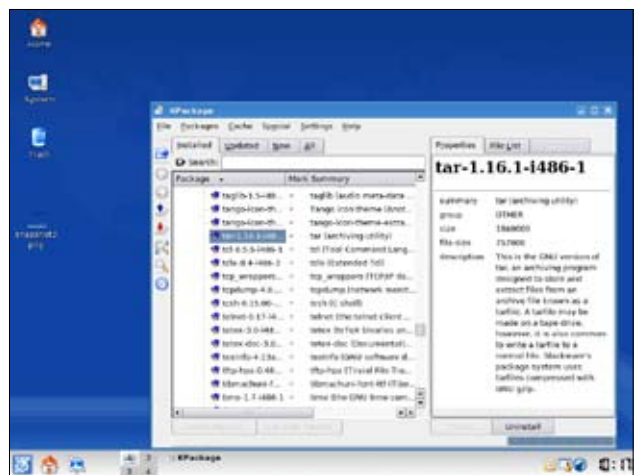


Figure 2: The KPackage utility

the *passwd* command.

This is the time to go back to the KDM screen—hit *Ctrl+Alt+F7*—select your desktop interface from the Menu button, where I chose KDE, and log in.

The Slack-perience

Slackware comes with KDE 3.5.10, which is the latest stable version from the KDE 3.x series. Figure 1 shows what the desktop looks like by default. This is actually the default KDE 3.5.10 desktop, including the wallpaper and the theme—Slackware hasn't even been customised to add a wallpaper of its own. Talk about conservation! If you are a GTK sort of a person, your option is XFCE 4.4.3—sorry folks, no GNOME here!

As for the other important software, it comes with kernel 2.6.27.7, GCC 4.2.4 and Xorg server 1.4.2. Since I opted for a full installation, I noticed that whether you're a normal desktop user, a developer, or a sys admin, Slackware doesn't disappoint anyone—things are pretty much covered for everyone here.

There's really nothing much to report here—everything is in line with vanilla upstream versions (check sidebar for a list of available GUI applications).

List of GUI applications	
Software Category	Available applications
Development	Kdevelop, Translation, Web Development, Cevisia (CVS fronted), KBugBuster (KDE bug management), KUIViewer, Kommander Editor, Umbrello (UML modeller), kjscmd (JavaScript console), etc.
Educational	Various language, mathematics, and science tools.
Games	Various arcade, board, card, tactics and strategy games.
Graphics	GIMP, GQview, KVDI, KFaxView, KGhostView, KSnapshot, KView, KolourPaint, Kooka (scanning and OCR program), Krita (image editor), KuickShow, etc.
Internet	Akregator, Firefox, Kget download manager, KMail, KNetAttach (network folder wizard), KNode news reader, Kandy (mobile phone tool), Pidgin, SeaMonkey, Thunderbird, Xchat, etc.
Mutimedia	Amarok, Audacious, JuK, K3b, KAudioCreator, KRec (recording utility), KsCD (CD player), Xine, etc.
Office	KOffice, KAlarm (alarm scheduler), KNotes, KOrganizer, KPlato (Palm Pilot tool), Kexi (database creator), etc.

Table 1

However, don't automatically take that conclusion as a bad sign. In fact, I meant it as a compliment. I found the desktop to be more stable than anything I've used in the last few years. Another thing to note is how snappy the desktop appears—it felt so much faster than the alternate distros, that I wonder why the others can't follow suit? Why do they have to be so slow?

I'd like to talk about one area, in particular—package management. Slackware doesn't come with a graphical package manager, so you would have to depend on KDE's Kpackage. This can be used to add, remove and update the packages. Take a look at Figure 2. As for a command line alternative, Slackware does have its own tools here: *installpkg* for installing, *removepkg* to remove packages, and *upgradepkg* to upgrade installed packages. In fact, there are two other commands that you might find useful: *explodepkg* to extract files without installing them, and *makepkg* to create a Slackware package from source files. A handy resource to check whether a Slackware package is already available for a task you need to perform is packages.slackware.it.

Although the distro provides you with a lot of utilities, I guess

people would rather prefer working on the programs they are familiar with. Similarly, I did not find certain software that I frequently use—VLC, Flash, and system monitor (Gkrellm) were missing. So, a quick Google search gave me the link to the Slackware 12.2 repository at repository.slack.eu. From here, you can download the packages that are not currently installed. If it's not here, I recommend checking out the individual software vendor's website first. If a Slack build is not available here, which is likely, Google for a Slackware build.

I downloaded VLC player directly from the VideoLAN website at www.videolan.org/vlc. You can install the VLC package in Slackware by simply using the following command:


```
installpkg vlc-0.9.8a-i486-2alien.tgz
```

Similarly, I downloaded the native Flash tar file provided by Adobe, and used the Adobe default installer to install it after going through the readme files.

After having installed these packages, I tried playing songs in my system. I noticed that I was not able to play music simultaneously on two media players. The reason was

that ALSA was not being configured out of the box, and this sound was default to the age-old OSS drivers. So I configured ALSA by running *alsaconf* from the command line, and that fixed the issue for me.

Another interesting thing is that Slackware comes with the KOffice suite for word processing, spreadsheet, presentation, or database management requirements. I didn't mind using it for a change—in fact, I am writing this article on KWord. However, if you can't do without OpenOffice.org, you can take a look at the repository hosted at rworkman.net/pkgs. In fact, the site has Slack builds for a lot of other useful utilities also.

In a nutshell, I found Slackware 12.2 a very stable version, very true to its motto. Whether I'll stick with it for long, I can't really tell—remember, I confessed that I'm fickle. But if you are someone who vouches for stability, then I've got to say Slackware is the distro for you. By the way, Debian 5 has also been released. ;-) 

Resources

- Slackware book: www.slackbook.org/html/index.html
- Third-party packages: rworkman.net/pkgs

Slackware 12.2



Pros:

Stable, less resource hungry compared to others, comes with more or less the latest apps.

Cons:

OpenOffice.org and GNOME missing, newbie unfriendly

Platform: x86 or x86-64

Price: Free (as in beer)

Website: www.slackware.com

By: Abhijit Paul Choudhury

The author loves to hack on open source and is a gamer by heart. Oh, and he's part of the LFY bureau too.

Electronic Components Products Machinery Will now be sold here

Presenting Electronics Bazaar,
India's First B2B Sourcing Guide
For Electronics Components,
Products and Machinery.

Unit Price: **Rs 50**

Subscription Rate: **Rs 350 (12 Issues)**

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area,
Phase-I, New Delhi 110 020; Ph: 011-26810601-03; Fax: 26817563;
Email: info@efyindia.com; Website: eleb2b.com



Infrastructure Box

Plugin Backup

Consulting

Neuron IT Solutions

Statement about ownership and other particulars about

LINUX FOR YOU

FORM IV (See Rule 8)

- | | | |
|--|---|---|
| 1. Place of publication | : | New Delhi |
| 2. Periodicity of its publication | : | Monthly |
| 3. Printer's Name | : | Ramesh Chopra |
| Nationality | : | Indian |
| Address | : | LINUX FOR YOU
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020 |
| 4. Publisher's Name | : | Same as (3) above |
| Nationality
and address | : | |
| 5. Names and addresses of
individuals who own the
newspaper & partners or
shareholders holding more
than 1% of the total capital | : | EFY Enterprises Pvt Ltd
D-87/1, Okhla Industrial Area,
Phase 1, New Delhi 110020 |

I, Ramesh Chopra, hereby declare that the particulars given above are true to the best of
my knowledge and belief.

Date: 28-2-2009

Ramesh Chopra
Publisher



Open Source

A Panacea for the Recession

As the recession grips more economies and enterprises, it's the perfect time to adopt the open source business model. We explain why.

*T*he technology landscape has been undergoing a massive transformation over the past couple of years. In the early 2000s, organisations were going after products and appliances irrationally, investing huge sums without any guarantees of a quantifiable ROI (return on investment). Today, awareness of technology has grown and organisations tend to demand a measurable ROI before adopting a new technology. Open source software has always been a good option to provide effective solutions within relatively low budgets. Not surprisingly, a number of organisations have started adopting open source software solutions. This trend has been boosted by the ongoing recession, with companies looking to cut costs.

Open source can mean big bucks!

The revenues of Red Hat grew by 14 per cent during the last economic bust in 2001-2002, which then increased to 38 per cent and 58 per cent in 2003 and 2004, respectively, demonstrating the increased usage of open source software during the economic crunch. Novell also showed interest in playing a significant role in the

open source software market in November 2003 by acquiring SUSE Linux for \$210 million. Last year, Sun Microsystems acquired MySQL AB, the developer of the world's most popular and fastest growing database, for \$1 billion. Writing open source software enables a company to gain access to the open community, which in turn helps in accelerating the pace at which an idea matures.

John Roberts, CEO and founder, SugarCRM, started his company in 2004 with a commercial open source concept. In one of his interviews given to Sramana Mitra, [www.sramanamitra.com/2008/12/11/5178/] in December 2008, he said, "I convinced two strong engineers at E.piphany to join me. We all resigned together and started SugarCRM without any angel or VC (venture capitalist) money. It was the three of us, each in his house with headphones on, writing and designing code and posting it up on SourceForge.net. We did that for three months. And soon enough, people all over the world started downloading the code." Early this year, SugarCRM joined hands with Tata Communications to make its software accessible to the Indian industry.

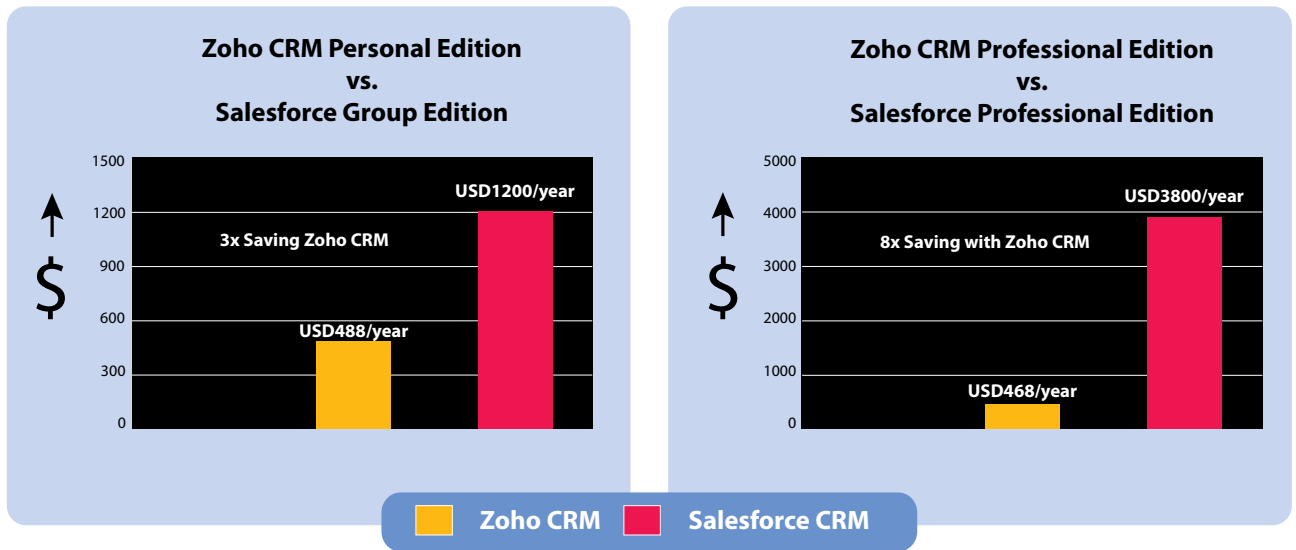


Figure 1: A comparison chart of Zoho services with those offered by its competitor, Salesforce

A matter of generating revenues


But how do open source software companies make money? There are a variety of business models to generate revenues from open source software:

- Releasing commercial extensions/plugin-ins to open source software
- Offering free community-based editions and paid commercial editions with more functionality and features
- Using free and open source to gain media attention, and attract users who might become potential customers for other commercial products
- Offering paid technical support along with free community-based support
- Making the software available via the Internet like on-demand applications, and offering paid subscriptions for online accounts and services

The point about paid technical support is particularly relevant in these times. A recent survey by IDC, a global market intelligence firm, suggests: “The economic slowdown in the United States may actually boost demand for open source services. If organisations adopt more open source software as part of a strategy to reduce software costs, the demand for related services should increase.” Many service providers have switched from expensive proprietary system management software to open source software like OpenNMS, Zenoss, Hyperic, Groundworks, Nagios, etc, so that they can cut down the technology cost overheads of their customers.

Finally, the SaaS (Software as a Service) business model has proved successful. SaaS delivers non-intrusive and hassle-free cloud-based products via a subscription-based model with no one-time costs involved—an organisation can pay as per software usage. The key to success for a SaaS player is obviously best technology usage, application features, and more

importantly, an optimised pricing model. Coming up with a competitive pricing model becomes very difficult if the technology is built over expensive hardware or proprietary software licences. This challenge provides a great opportunity to utilise open source software as SaaS-ready solutions. The success enjoyed by companies like SugarCRM, Zoho (a Chennai-based offshoot of AdventNet) and Zimbra (acquired by Yahoo! in 2007 for \$350 million), follows the same concept—enabling organisations to reduce their IT expenditure and increase flexibility by leveraging open source software. Zoho prepared a cost saving chart that compared its services with those offered by its competitor, Salesforce, highlighting the difference in costs between the two (see Figure 1).

According to Springboard Research, the Indian SaaS market is set to reach \$165 million by 2010, due to a compound annual growth rate (CAGR) of 77 per cent from 2006 to 2010. We might see greater demand for software available under SaaS model running over private clouds or corporate networks than public clouds due to data security and compliance regulations, but small and medium-sized companies can utilise the software over public clouds to boost their performance at a much lower cost. The increasing demands for such models provides a huge opportunity for service providers to make a shift from traditional delivery models to avant-garde technology models. The economy might be having a tough time, but if they play their cards right, open source companies might never have had it so good. **END** 

By: Dhruv Soi

The author is the founder and principal consultant, Torrid Networks, and chair, OWASP India. He can be reached at: dhruv.soi@torridnet.com

A Matter of Recession



“You seem a bit lost. Is this the first time you have visited our firm?”

“Yes, sir. I...”

“It is an amazing place. We started out from a room in a hostel a couple of years ago and today we are a million-dollar organisation...and still growing. In all directions!”

“Yes, sir. I was...”

“Looking around? Yes, do so by all means. You can also visit the demo labs and try out some of the products we have come out with in recent times. As well as some of the old ones, actually -- though there are some OS compatibility issues. You could try out some of our “crash, burn and die in seven minutes” betas too. They are a new concept in demos—if you use them for too long, they format your hard drive and set fire to your keyboard!”

“Well, I was looking for...”

“We have tie-ups with some of the biggest firms in the business. We operate in six countries and offer round-the-clock support for all our products, 60 seconds a minute, 60 minutes an hour, 24 hours a day, seven days, two weeks a fortnight, two fortnights a month, twelve months a year and ten years a decade. After that we will see! And we are proactive too—when our customers did not upgrade their products, we encouraged them to do so by telling them that their existing copies would expire in two days, unless upgraded.”

“Yes, I do know that. The famous “Upgrade or crash. Karo ya maro” campaign, wasn’t it? I...”

“So glad you have seen the campaign. It got us so much visibility. Some people were annoyed at having to shell out extra cash for the upgrade, but then, we cannot let costs come in the way of progress, can we?”

“It was very innovative...”

“Well, one of the reasons for success has been the fact that we have been able to offer solutions that are low-cost in comparison to their counterparts. We like to stay one step ahead of the competition—sometimes even two steps. To innovate constantly, we breed a culture of innovation. From the CEO who uses three different phones to write one SMS; to the newspaper boy who is encouraged to wrap the newspapers in new shapes every day; to the CEO’s driver who finds new ways to park his car in the parking spaces of other companies; to ...”

“Actually, sir, what I was looking for...”

“A lot of people have been saying that the current recession is bad for business. Nonsense, I say. It is in fact a tremendous opportunity for us. As people are looking to cut costs, we can come up with more economical solutions. We have even come out with a Pink Slipper—a device that can be used as footwear and also to generate dismissal letters by using a slide out QWERTY keypad that can be activated by moving both straps seventeen degrees to the right.”

“The recession...”

“A great opportunity for all innovators. We can redefine businesses, focus on core competencies, and best of all, get access to talent at reasonable costs...”

“Sir, that is what I wanted to talk about.”

“You are an HR consultant and want to show us your database of potential employees?”

“No, sir. I...”

“Ah, then you wish to invest in our company as our share prices have declined a bit. An excellent idea! I know our prices have dipped and that we currently have to pay you to buy our shares but that will change. In our world, change is the only constant. Let me forward you to...”


“No, sir. I am applying for a job. Our company is slashing jobs because of the recession...”

“They are? Do you know if they are using Pink Slipper?”

“Sir, please, I just want to submit my resume...”

“Well, you can, my dear chap, but we are not hiring at the moment. The recession, you know...”

“But, sir, you were just talking about how the recession was a great opportunity...”

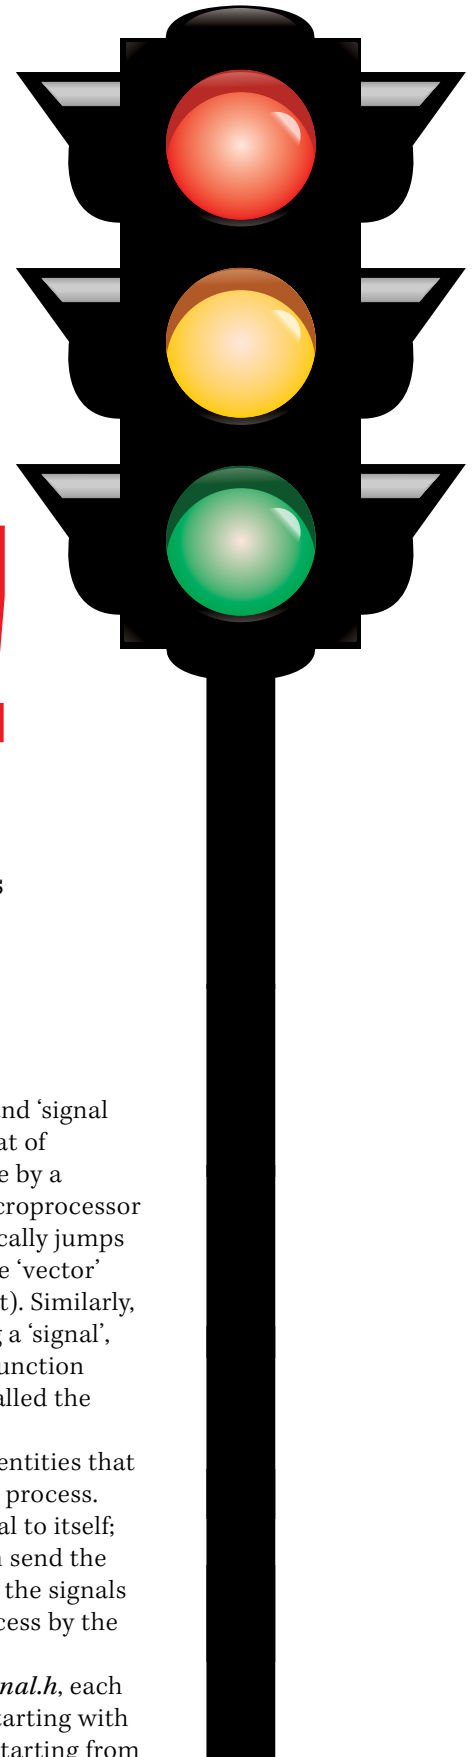
“So it is. Right now, it is giving us the opportunity to turn down your application. All the very best. Do visit us again.” 

Nimish Dubey

The author is a writing practitioner who believes that laughter is the best medicine, especially if the dosage includes PG Wodehouse, Stephen Leacock, Spike Milligan and Tom Sharpe. You can reach him at nimishdubey@gmail.com

This feature is a reprint and was first published in the Jan '09 edition of 'i.t.' magazine, a sister publication of LFY.

Watch Out for the Signals!



What in the world is the 'signals' framework and how can systems programmers make use of it?

*W*e say people are clever when they understand the 'signals' in real life. The case is the same with Linux too. The right signals sent across at the right time in the system make it fast and responsive. This article throws light on this 'signals' framework in the Linux system, and explores how systems programmers can make use of it.

The framework

The motivation behind this framework of 'signals' is to make the process aware that something has happened in the system, and the target process should perform some predefined set of actions to keep the system running smoothly. These actions range from 'self-termination' to 'clean-up'.

The concept of 'signals' and 'signal handling' is analogous to that of the 'interrupt' handling done by a microprocessor. When a microprocessor receives an interrupt, it typically jumps to a fixed location (called the 'vector' location for a given interrupt). Similarly, in Linux, a process receiving a 'signal', typically invokes a specific function registered with the signal, called the 'handler' for a given signal.

There could be multiple entities that can send a 'signal' to a given process. The process can send a signal to itself; other running processes can send the signals to a given process or the signals could be sent to a given process by the Linux kernel too.

As defined in *include/signal.h*, each signal has a specific name starting with 'SIG' and a unique number starting from

1. Each signal also has a default action associated with it. It could be either of the four mentioned below:

- **Exit:** Makes the process exit.
- **Core:** Forces the process to exit and create a core dump file.
- **Stop:** Stops/suspends the process.
- **Ignore:** No action taken.

The framework is also flexible enough, so you can change the default disposition of a signal to one of your choice by overriding the default signal handler.

The framework also allows a given process to block some signals so that they don't get delivered to the process at all.

One restriction imposed here by the system is that the process cannot change the default disposition for **SIGKILL** and **SIGSTOP**.

Practical applications

This section talks about the various standard applications based on 'signals' in a typical Linux system.

IPC 302: Yes! This is the basic one, or the **kill** command. It sends the **SIGTERM** to a process specified in the argument, and the process terminates. You can also send **SIGKILL** to kill a process through the same command if the process ignores **SIGTERM**.

Ctrl+C: When you press **Ctrl+C** on the keyboard, the process running on the foreground on the given terminal receives **SIGINT**. The process also terminates by default when it receives **SIGINT**.

Old buddy GDB: Yes, the working of GDB is totally based on the signals **SIGSTOP** and **SIGCONT**. When the process being debugged reaches a breakpoint, GDB sends **SIGSTOP** to a given process and its execution halts. Now, after looking at the available information, when a user makes the process run, GDB sends out **SIGCONT** which clears **SIGSTOP** and lets the process go ahead.

Alarms: When an application wants to run the timers, they typically make use of APIs like **setitimer()** specifying the time out value in the arguments. When the timer expires, the process gets **SIGALARM**.

Child care: In the Linux system, process creations are done through the **fork()** system call and the processes have a parent-child relationship. Now

parents need to be informed when a child changes its state so that they can take appropriate action, such as doing a cleanup, spawning one more child if one gets killed, etc. This functionality is achieved through **SIGCHLD**.

Broken pipe: In Linux, processes pass the data to other processes through pipes/fifos/sockets, etc. When a process attempts to write to a broken pipe, the process receives **SIGPIPE** indicating the same.

Check your Memory/Math/Instruction set: When a process attempts to access an invalid memory address, it receives a **SIGSEGV** from the system. Similarly, when a program attempts to execute invalid floating point computation, it receives **SIGFPE** for it. Also, if a process attempts to execute illegal instruction, it gets **SIGILL** indicating the same.

By default, these signals also result in programs to crash with a core dump.

Something left for programmers: There are two signals **SIGUSR1** and **SIGUSR2**, which are left for the programmers, and their meanings need to be set by them.

API support for signal handling

There are two main APIs available for programmers to change the default disposition of the signals. The first is **signal()**, which looks like what's shown below:

```
void (* signal(int sig, void (*func)(int)))(int);
```

The equivalent **typedefd** version for the same, which is easier to read, is as follows:

```
typedef void (*sig_t) (int);
```

```
sig_t signal(int sig, sig_t func);
```

The function is very simple to use. You only need to specify the signal number and call-back function that needs to be registered. But the API is getting deprecated, and a more robust and elaborate API called **sigaction()** is available:

```
int sigaction(int sig, const struct sigaction *act, struct sigaction *oact);
```



The *sigaction* structure includes the following members:

```
struct sigaction {
    void (*sa_handler)();
    void (*sa_sigaction)(int, siginfo_t *, void
*);
    sigset_t sa_mask;
    int sa_flags;
};
```

The function could be used to get and modify the disposition for a specified signal.

The *sigprocmask()* API is used to block/unblock a specified signal and *sigsuspend()* is used to suspend the process till it receives a specified signal or the process gets killed.

The *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()* are the auxiliary functions available and should be used to operate on *sigset_t*.

A piece of code to catch SIGINT

Here is a simple piece of code to catch *SIGINT*:

```
/****** SigInt.c
******/
#include <signal.h>
#include <stdio.h>
#include <unistd.h>

void SigIntHandler(int sig)
{
    printf("Received signal %d\n", sig);
}

int main()
{
    struct sigaction act;
    int count = 5;

    act.sa_handler = SigIntHandler;
    sigemptyset(&act.sa_mask);
    act.sa_flags = 0;

    sigaction(SIGINT, &act, 0);

    while(count-- > 0) {
        printf("Looping!\n");
        sleep(5);
    }
}
```

```
}
}
```

Here we compile and run the code:

```
# gcc SigInt.c
# ./a.out
Looping!
Looping!
Received signal 2
Looping!
Looping!
Received signal 2
Looping!
#
```


Watch for some pitfalls

The following are some points to be considered during the design time:

1. A process might wait for an indefinite period of time when it invokes *sigsuspend()* but does not receive the desired signal.
2. Due to a timing mismatch, a process might wait for a signal that has already occurred.
3. The child process inherits the same signal handlers from the parent after *fork()*.
4. The final and most important point is that the signal handlers should be short and reentrant.



A nice article about signals and reentrancy is available at www.ibm.com/developerworks/linux/library/l-reent.html

Here onwards

This article talks about the basics of the 'signals' framework on Linux. Some more advanced and interesting stuff is available in a *Linux Journal* article on signals: www.linuxjournal.com/article/3985. 

By: Nilesh Govande

The author is a Linux enthusiast and can be contacted at nileshgovande@yahoo.com. His areas of interest include Linux systems software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.

Open Source Solutions

Linux and Open Source Support


Training and Linux Staffing.

Mail Server

- Push Mail to hand held devices
- Mail Archival
- Customisation of Email to Workflow
- Active-Active Clustering
- Using central LDAP/AD authentication
- Coexistence with Exchange/Other Mail Servers
- Fully customisable thru GUI

Recent Project

Mail Server for more than 2000 users distributed across 40 locations in India with distributed servers for each location and the facility for any user to check mail from anywhere and facility to move user mailboxes in one click on change of location



Exam Centre for RHCE & RHCSS

India's First RHCSS*

<p>Keen & Able Computers Pvt Ltd 55, Deepak Building 13, Nehru Place New Delhi - 110 019 Tel : +91-11-26426184/5, +91-11-30880046-48 FAX : +91-11-41808565</p>	<p>FOSTERing Linux Free & Open Source Training Education & Research SCO M-22, II Floor Old DLF Colony Gurgaon - 122 001 Tel : +91-124-4268187, +91-9212800496</p>
---	--

Email: info@fosteringlinux.com
 Website: www.keenable.com



The Smart, Simple, Powerful IPC

Let's learn the intricacies of D-Bus and use it to hack some nifty features into programs.

*I*nter-process Communication (IPC) helps applications to talk to each other. You might have seen Firefox automatically tuned to offline mode when your Internet connection is down. Ever wondered how this happens? This is because the NetworkManager application talks to Firefox using a back-end utility called D-Bus to update it on the status of the Internet connection.

D-Bus (Desktop Bus) is a simple IPC, developed as part of freedesktop.org project. It provides an abstraction layer over various applications to expose their functionalities and possibilities. If you want to utilise some feature of an application to make another program perform a specific task, you can easily implement it by making the process D-Bus aware. Once an application is made D-Bus compliant there's no need to recompile or embed code in it to make it communicate with other applications.

One thing really cool about D-Bus is that it helps developers write code for any D-Bus compliant application in a language of their choice. Currently, D-Bus bindings are available for C/C++, Glib, Java, Python, Perl, Ruby, etc.

Understanding D-Bus

D-Bus is a service daemon that runs in the background. We use bus daemons to interact with applications and their functionalities. The bus daemon forwards and receives messages to and from applications. There are two types of bus daemons: SessionBus and SystemBus.

The daemon that is attached to each user session is called SessionBus. When a user logs in, applications launched by him are attached to the SessionBus—a local bus limited to communicating between desktop applications that belong to a specific user currently logged in.

On the contrary, SystemBus is system-wide. It is initiated when the system boots, and is 'global' to the operating system. It is capable of interacting with the kernel and various system-wide events. Hardware Abstraction Layer (HAL), NetworkManager and udev are applications that use SystemBus.

In this article, I will use Python bindings to explore the D-Bus daemon. To begin with, if we want to use a desktop-level conversation, a SessionBus object can be created as follows:

```
[slynux@slynux-laptop dbus-python-0.83.0]$ python  
Python 2.5.2 (r252:60911, Sep 30 2008, 15:41:38)  
[GCC 4.3.2 20080917 (Red Hat 4.3.2-4)] on linux2
```

Type "help", "copyright", "credits" or "license" for more information.

```
>>> import dbus
>>> bus = dbus.SessionBus()
>>>
```

While a `SystemBus`, on the other hand, can be created by simply replacing the `dbus.SessionBus()` element in the above code to `dbus.SystemBus()`:

```
>>> bus = dbus.SystemBus()
```

Every application that intends to share its objects and methods are started as D-Bus services. A D-Bus enabled application exports its objects with their functionalities as methods that other applications can use. By connecting to the corresponding bus and the application object, the application's functionalities can be accessed from other applications.

We use an addressing method to identify each application and its functionalities—reversed domain name addressing. For example, `NetworkManager` is addressed as `'org.freedesktop.NetworkManager'`, `Pidgin` as `'org.gnome.Pidgin'`, etc.

Each of the applications can export numerous objects and functions—that is, `NetworkManager` has got different parameters such as 'if network is up or down', 'the current active wifi profile', etc.

Proxy objects and interfaces

The term 'proxy objects' refers to objects that point to remote applications and are accessed through D-Bus session. Let's explore how to create proxy objects.

To obtain a proxy object, call the `get_object` method on the bus. For example, `NetworkManager` has the well-known name `org.freedesktop.NetworkManager` and exports an object whose object path is `/org/freedesktop/NetworkManager`, plus an object per network interface at object paths like `/org/freedesktop/NetworkManager/Devices/wlan0`.

```
>>> import dbus
>>> bus = dbus.SystemBus()
>>> proxy_object = bus.get_object('org.freedesktop.NetworkManager', '/org/freedesktop/NetworkManager')
```

The format of the parameters for `get_object()` is `get_object(dbus_service_name, object_path)`. So, you can see from the above code snippet, `org.freedesktop.NetworkManager` is the service name and `/org/freedesktop/NetworkManager` is the object path. The object path is different for accessing different objects specified by the service. Here a proxy object referring to the `NetworkManager` is created. Now it is possible to access different properties of this object. For example, we can check whether the `NetworkManager` is in sleep or wake mode, or if it is connected to some network or not, as follows:

```
>>> print proxy_object.state() # To know the NM state
```

4

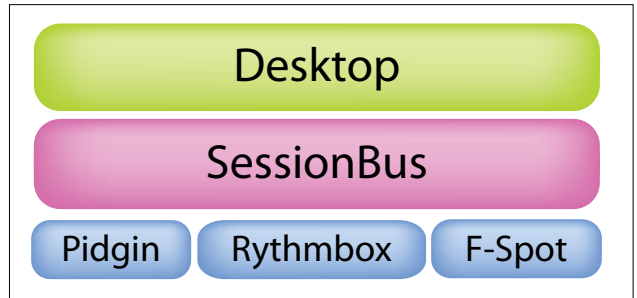


Figure 1: D-Bus SessionBus

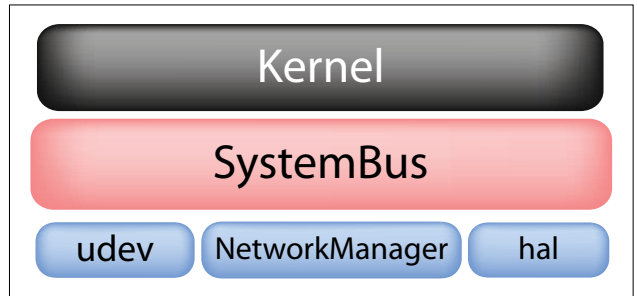


Figure 2: D-Bus SystemBus

The returned integer in the above example is called the `NM_STATE`. This corresponds to following states:

- `'NM_STATE_UNKNOWN = 0'` means the `NetworkManager` daemon is in an unknown state.
- `'NM_STATE_ASLEEP = 1'` means the `NetworkManager` daemon is asleep and all interfaces managed by it are inactive.
- `'NM_STATE_CONNECTING = 2'` means the `NetworkManager` daemon is connecting to a device.
- `'NM_STATE_CONNECTED = 3'` means the `NetworkManager` daemon is connected.
- `'NM_STATE_DISCONNECTED = 4'` means the `NetworkManager` daemon is disconnected.

Let's take a look at the following code:

```
>>> proxy_object.sleep() # Disable NetworkManager
>>> proxy_object.wake() # Enable NetworkManager
>>> proxy_object.GetDevices()
dbus.Array([dbus.ObjectPath('/org/freedesktop/Hal/devices/net_00_1c_23_fb_37_22'), dbus.ObjectPath('/org/freedesktop/Hal/devices/net_00_1c_bf_87_25_d2')], signature=dbus.Signature('o'))
```

You can see that the code lists objects of two network interfaces with MAC ID `00:1c:bf:87:25:d2` and `00:1c:23:fb:37:22` along with their HAL object paths. The `dbus.Array` element is a D-Bus object specific data type. We'll discuss more on D-Bus types in the later part of the article.

An object path can support any number of different interfaces. Before calling any method, you need to specify which interface you want to use. Interfaces are sub-objects that can be used to refer to a group of other objects to provide a higher level of abstraction on proxy objects and their exported methods. It provides a name-spacing mechanism. You can have a better understanding of the concepts of



Figure 3: Different interfaces provided by same object

Interfaces from Figure 3.

Take a look at the following code:

```
>>> bus=dbus.SystemBus()
>>> proxy_object=bus.get_object('org.freedesktop.NetworkManager',
'/org/freedesktop/Hal/devices/net_00_1c_bf_87_25_d2')
>>> proxy_object.GetAccessPoints(dbus_interface='org.freedesktop.
NetworkManager.Device.Wireless')
dbus.Array([dbus.ObjectPath('/org/freedesktop/NetworkManager/
AccessPoint/4'), signature=dbus.Signature('o')])
# Above method returns with a dbus Array type containing object path
of currently available access points. dbus_interface='org.freedesktop.
NetworkManager.Device.Wireless'

>>> proxy_object.Get('/org/freedesktop/Hal/devices/net_00_1c_23_fb_37_
22','HwAddress',dbus_interface='org.freedesktop.dbus.Properties')
dbus.String(u'00:1C:23:FB:37:22', variant_level=1)
# It returns Hardware Address of the Interface. dbus_interface='org.freedesktop.
dbus.Properties')
```

Here we have used two different interfaces under the same object path. The D-Bus bindings provide an object type of *Dbus.Interface*, making it easier to interpret. We can rewrite the above code as follows:

```
>>> hw_address_interface = dbus.Interface(proxy_object,dbus_interface='org.
freedesktop.dbus.Properties')
>>> hw_address_interface.Get('/org/freedesktop/Hal/devices/net_00_1c_23_
fb_37_22','HwAddress')
```

Even though both are same, the latter eliminates the need for specifying the interfaces parameter *dbus_interface* every time we call a method.

The D-Bus package comes with a set of utilities to manage the D-Bus daemon activities. The *dbus-monitor* is one such utility that is used to keep track of all active D-Bus sessions in a running system. It helps you be aware of the applications that make use of D-Bus and its events:

```
[slynux@slynux-laptop ~]$ dbus-monitor
signal sender=org.freedesktop.dbus -> dest=:1.134 path=/org/freedesktop/
```

```
dbus; interface=org.freedesktop.dbus; member=NameAcquired
string ":1.134"
method call sender=:1.134 -> dest=org.freedesktop.dbus path=/org/
freedesktop/dbus; interface=org.freedesktop.dbus; member=AddMatch
string "type='method_call'"
method call sender=:1.134 -> dest=org.freedesktop.dbus path=/org/
freedesktop/dbus; interface=org.freedesktop.dbus; member=AddMatch
string "type='error'"
signal sender=:1.54 -> dest=(null destination) path=/im/pidgin/
purple/PurpleObject; interface=im.pidgin.purple.PurpleInterface;
member=BuddyIconChanged
int32 24422
signal sender=:1.54 -> dest=(null destination) path=/im/pidgin/
purple/PurpleObject; interface=im.pidgin.purple.PurpleInterface;
member=DrawingTooltip
```

The utility gives you an overview of the different D-Bus events and the applications using D-Bus. As you can see in the output, an event related to Pidgin 'BuddyIconChanged' along with some other D-Bus events has taken place.

dbus-launch and *dbus-sendto* are two other utilities available for working with D-Bus. Check out their man pages to understand the purpose of these utilities. *dbus-sendto* can be used to interact with the buses and their return strings. It can be used if we want to write pure Bash-coded applications.

D-Bus activation

We can start a D-Bus service such as *org.gnome.example_service* from a server program or we can start a service by calling it by name. The technique of starting a service by name is called D-Bus activation. There are several instances where we need to start another application to make some feature of the currently running application work. For example, consider a video editor which extracts still images from the GNOME Web cam tool Cheese. Since the video editor needs Cheese to be running, it needs to be started. If Cheese is defined as a D-Bus service, we can easily start Cheese by D-Bus activation.

Most of the applications which make use of D-Bus are defined as D-Bus services. You can have a look at the contents of the */usr/share/dbus-1/* directory for the some available services:

```
[slynux@slynux-laptop services]$ ls /usr/share/dbus-1/services/ | tail
org.gnome.keyring.service
org.gnome.PolicyKit.AuthorizationManager.service
org.gnome.PolicyKit.service
org.gnome.Rhythmbox.service
org.gnome.SettingsDaemon.service
org.gnome.Tomboy.service
org.gtk.Private.GPhoto2VolumeMonitor.service
org.gtk.Private.HalVolumeMonitor.service
org.xchat.service.service
sealart.service
```

Each of these services can be started by using the *start_service_by_name()* method.

For example, the Tomboy note-taking application can be launched by running the following from a Python shell:

```
>>> import dbus
>>> bus=dbus.SessionBus()
>>> bus.start_service_by_name('org.gnome.Tomboy')
(True, dbus.UInt32(1L))
```

You can see that Tomboy is started and the function returns True.

In fact, it is very easy to create D-Bus services. Create a text file, called *org.gnome.Newservice.service* for example, with following contents:

```
[D-BUS Service]
Name=org.gnome.Newservice
Exec=/usr/bin/newservice
```

Now you can start Newservice by name.

Data types and type casting

Since D-Bus is an inter-process message passing mechanism, it deals with various data types, depending on the data to be received or sent. One of the primary benefits of D-Bus is that it is flexible with data type conversions. Since we are more concerned with D-Bus in Python's context, let us take a look at how D-Bus types and Python types are tuned to each other with auto typecasting. D-Bus uses static types. Since Python types and D-Bus types are compatible to each other, we never have to worry about type conversion hurdles.

Table 1 lists the types supported and their conversions. Types marked (*) may be a subclass of either *int* or *long*, depending on the platform.

From the above table, you can infer that if we have some string to be passed or received through D-Bus daemon, it is received or sent as its equivalent D-Bus type. Likewise, *string* is send as *dbus.String("string")*.

We can call methods provided by the proxy object in two ways—synchronous call or asynchronous call. Synchronous calls block any other methods to be called until the current function call ends and returns something. Asynchronous (non-blocking) method calls allow multiple method calls to be in progress simultaneously, and allow your applications to do other work while it waits for results/answers. Asynchronous calls are invoked by setting up an event loop like *Gmainloop* or *gtk.main()*.

Hands-on D-Bus client-server

Let us code a simple ExampleObject to be exported under the *org.example.Sample* service and a client application, to understand programming with D-Bus better:

- D-Bus service: *org.example.Sample*
- File name: *dbus-example-service.py*

Python type	Converted to D-Bus type	Notes
D-Bus proxy object	ObjectPath (signature 'o')	(+)
dbus.Interface	ObjectPath (signature 'o')	(+)
dbus.service.Object	ObjectPath (signature 'o')	(+)
dbus.Boolean	Boolean (signature 'b')	a subclass of int
dbus.Byte	byte (signature 'y')	a subclass of int
dbus.Int16	16-bit signed integer ('n')	a subclass of int
dbus.Int32	32-bit signed integer ('i')	a subclass of int
dbus.Int64	64-bit signed integer ('x')	(*)
dbus.UInt16	16-bit unsigned integer ('q')	a subclass of int
dbus.UInt32	32-bit unsigned integer ('u')	(*)_
dbus.UInt64	64-bit unsigned integer ('t')	(*)_
dbus.Double	double-precision float ('d')	a subclass of float
dbus.ObjectPath	object path ('o')	a subclass of str
dbus.Signature	signature ('g')	a subclass of str
dbus.String	string ('s')	a subclass of unicode
dbus.UTF8String	string ('s')	a subclass of str
bool	Boolean ('b')	
int or subclass	32-bit signed integer ('i')	
long or subclass	64-bit signed integer ('x')	
float or subclass	double-precision float ('d')	
str or subclass	string ('s')	must be valid UTF-8
unicode or sub-class	string ('s')	

Table 1

```
#!/usr/bin/env python
import gobject
import dbus
import dbus.service
import dbus.mainloop.glib
```

```
class ExampleObject(dbus.service.Object):
    @dbus.service.method("org.example.Sample",
        in_signature='s', out_signature='as')
    def HelloWorld(self, test_message):
        print (str(test_message))

        return ["Hello World ", dbus.service.str(test_message)]
```

```
@dbus.service.method("org.example.Sample",
    in_signature="", out_signature='s')
def Ping(self):
    print "Pinged"
    return str("Hi. I am Alive")

@dbus.service.method("org.example.Sample",
```



Figure 4: ExampleObject and available methods

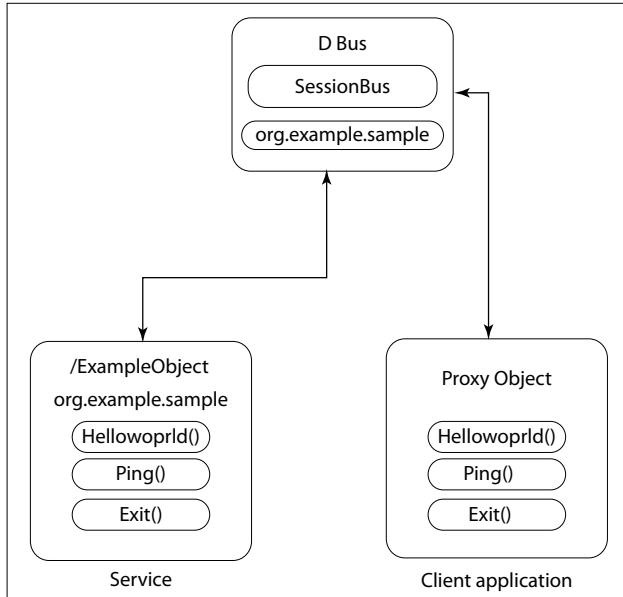


Figure 5: A schematic representation of service-client interaction

```

in_signature='', out_signature='')

def Exit(self):
    mainloop.quit()

if __name__ == '__main__':
    dbus.mainloop.glib.DBusGMainLoop(set_as_default=True)

    session_bus = dbus.SessionBus()
    name = dbus.service.BusName("org.example.Sample", session_bus)
    object = ExampleObject(session_bus, '/ExampleObject')

    mainloop = gobject.MainLoop()
    print "Running example dbus service: org.example.Sample"
    mainloop.run()

```

Here we have a class derived from *dbus.service.Object*, which consists of the *HelloWorld()*, *Ping()* and *Exit* functions that are to be exposed through the service. The decorator like *@dbus.service.method("org.example.Sample", in_signature="", out_signature="")* is used to expose these functions. It consists of parameters *in_signature* and *out_signature*, specifying the type of input (parameters to the function) and output (return type). You can refer to Table 1 for the types that are available. For example, 's' specifies string, 'sa' specifies string array, 'i' specifies integer and so on.

Let us now code a D-Bus client (*client.py*) to access methods exported by *org.example.Sample*:

```
#!/usr/bin/env python
```

```

import dbus

bus = dbus.SessionBus()
remote_object = bus.get_object("org.example.Sample", "/ExampleObject")
interface = dbus.Interface(remote_object, 'org.example.Sample')
reply = interface.Ping()

print "Ping() returns : " + reply
reply = interface.HelloWorld("GNU/Linux")

```

```

print "Helloworld() returns: "
for s in reply:
    print s,

```

If you go through the above code, you can understand that it simply creates a proxy object and an interface to the *org.example.Sample* service. Further, it calls the methods available. You can call it through any type of D-Bus client access method like *dbus-send* tool. Try this:

```

[slynux@slynux-laptop examples]$ dbus-send --session \
--dest=org.example.Sample --print-reply \
/ExampleObject org.example.Sample.Ping
method return sender=:1.326 -> dest=:1.364 reply_serial=2
string "Hi. I am Alive"

```

Now, you can open a terminal and execute the service script first and client after that.

On terminal tab 1:

```

[slynux@slynux-laptop examples]$ python example-service.py
Running example dbus service: org.example.Sample
Pinged
GNU/Linux

```

One terminal tab 2:

```

[slynux@slynux-laptop examples]$ python example-client.py
Ping() returns : Hi. I am Alive
Helloworld() returns:
Hello World  dbus-service GNU/Linux

```

Hacking other applications with D-Bus

Let us now focus more on the implementation and go through the coding part involving some applications, say, for example, Pidgin. Pidgin is a well-known IM client that a lot of us use to talk to people. We will now work on the D-Bus service interfacing with Pidgin in order to talk with Pidgin:

```

#!/usr/bin/env python
import dbus, subprocess, time

def set_status(message):
    current = purple.PurpleSavedstatusGetType(purple.PurpleSavedstatusGetCurrent())
    status = purple.PurpleSavedstatusNew("", current)
    purple.PurpleSavedstatusSetMessage(status, message)
    purple.PurpleSavedstatusActivate(status)

```



```
bus = dbus.SessionBus()
obj = bus.get_object('im.pidgin.purple.PurpleService', '/im/pidgin/purple/
PurpleObject')
purple = dbus.Interface(obj, 'im.pidgin.purple.PurpleInterface')
while True:
    fortune=subprocess.Popen('fortune', stdout=subprocess.PIPE).
stdout.read()
    set_status(fortune)
    time.sleep(10)
```

The above script makes use of the *fortune* command to generate random quotes. You may have noticed the gnome-panel applet Fish. Do you remember the “free the fish” Easter egg? Fish uses *fortune* as its back-end for generating quotes. The above script sets the status message for Pidgin every 10 seconds with a random quote generated by the *fortune* command.

The next application in line is Tomboy, a note-taking application, which ships with GNOME. This is how you can talk to Tomboy and collect all the notes created with it to print them on a terminal:

```
#!/usr/bin/env python

import dbus

bus = dbus.SessionBus()

obj = bus.get_object('org.gnome.Tomboy', '/org/gnome/Tomboy/
RemoteControl')
tomboy = dbus.Interface(obj, 'org.gnome.Tomboy.RemoteControl')

notes = tomboy.ListAllNotes();

for note in notes:
    print tomboy.GetNoteContents(note)
```

How about fiddling with Exaile music player, the GNOME-based Amarok clone? Our aim is to write few lines of Bash script to enquire the application on current music track, album name and artist. Add the following lines to the *~/.bashrc* file:

```
artist=$(dbus-send --print-reply --dest=org.exaile.dbusInterface \
/dbusInterfaceObject org.exaile.dbusInterface.get_artist 2> /dev/null | grep '".*"'
-o | tr -d '"');

album=$(dbus-send --print-reply --dest=org.exaile.dbusInterface \
/dbusInterfaceObject org.exaile.dbusInterface.get_album 2> /dev/null | grep
'".*"' -o | tr -d '"');

if [[ -n $album ]]; then

echo -e "\nCurrently Playing $album, $artist\n";
fi
```

Notice how every time you open a new terminal, it lists the information about the song currently playing in Exaile. Of course, if the player is not running, it won't print anything. Here, the *dbus-send* command is used to communicate with Exaile through the D-Bus interface.

Finally, let's hack GNOME's PowerManager to hibernate our machine:

```
[slynux@slynux-laptop ~]$ python
Python 2.5.2 (r252:60911, Sep 30 2008, 15:41:38)
[GCC 4.3.2 20080917 (Red Hat 4.3.2-4)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>>
>>> import dbus
>>> bus=dbus.SessionBus()
>>> power = bus.get_object('org.freedesktop.PowerManagement', '/org/
freedesktop/PowerManagement')
>>> pm = dbus.Interface(power, 'org.freedesktop.PowerManagement')
>>> pm.Hibernate()
```

The above code makes the PowerManagement daemon execute the *Hibernate()* function and the machine goes into hibernation. You can also use the *Shutdown()*, *Reboot()*, *Suspend()* instead.

So you see, by embedding any kind of D-Bus interfacing you are able to extract different sorts of things from an application. There are numerous applications that are hackable with D-Bus interfacing. Try it out for yourself—it's fun!



Note: Debugging D-Bus applications can be a hurdle sometimes. You can use the *dbus-monitor* to examine the events for a better understanding. Alternatively, you can also check out the D-Feet D-Bus debugger tool written by John Palmeri.

Bottom line

Now a days, most of the GNOME and KDE apps come with dbus interface support. This makes it easier for applications to communicate with each other and eliminates the higher-degree task of recompiling every application to make it compatible with another.

Now, here is your task. You may find that some of your favourite applications do not have D-Bus support. If you do, maybe you can start writing the D-Bus interfacing for your favourite applications—contribute back to the community, it's not that hard really! Happy Hacking!



By: Sarath Lakshman

The author is a Hactivist of Free and Open Source Software from Kerala. He loves working on the GNU/Linux environment and contributes to the PiTiVi video editor project. He is also the developer of SLYNIX, a distro for newbies. He blogs at www.sarathlakshman.info

Programming in Python for Friends and Relations, Part 11



Secure Communication

Here's a simple application to help us stop making silly mistakes while communicating over e-mail.

The faint smile on your face turns into an expression of panic the moment you check your e-mail. “Oh, my God! The CEO has just forwarded an e-mail to Kiran, a union leader, instead of Kiran, the CFO.” But you manage to save the day. The recipient is on leave, and you delete the mail from his inward queue.

Your HR chief asks you to check your e-mail. When you do so, you can't help but smile. The CEO has just declared a holiday for his birthday! But the smile doesn't last long. The HR chief asks you to find out who sent that e-mail, as it surely wasn't the CEO. You are reminded of the line routinely printed by banks on their statements: “This is a computer-generated statement and does not require a signature!”

Increasingly, your financial dealings are online.

The statements are being sent by e-mail. To minimise the chances of the wrong person viewing confidential information, the statements are password protected. However, the passwords aren't very strong. They protect against casual snooping, which is fine for most of us. But you do need to put some effort into figuring out the password for each statement. It's not easy for you and your colleagues to do so.

There has to be a better way to manage all these scenarios.

Public key infrastructure

Public key-based algorithms have been around for about as long as I have been in the software field. Ubuntu owes its existence to the money made from the sale of Thawte, which issues digital certificates.

PGP (Pretty Good Privacy) came into existence in the early 1990s and GPG (GNU Privacy Guard) that conformed to the OpenPGP standard was available by the end of the 90s.

I have used a public key only to start an SSH session on a remote computer without having to give a password. But I have relied on GPG whenever I installed packages from a Fedora repository. Keeping the private key safe is a critical part of these security processes. The fear that the signing key may have been compromised resulted in the closure of the Fedora repositories for a noticeable period of time.

One reason for the lack of applications using OpenPGP may be that it is hard to get started with them. It is important to realise that this technique is based on people trusting each other and not on a third-party certificate. Would I trust the keys more if the issuing company had been audited by, say, PWC? A transaction between two parties does not need a certificate from a third party.

Before getting into programming using GPG, let us consider the steps involved in using the public key infrastructure with an e-mail client, Evolution. We choose Evolution as it comes with GPG support. Many e-mail clients now support OpenPGP -- for example, Sylpheed. Thunderbird requires the Enigmail plug-in, which, unfortunately, was not compatible with the x64 application I was using. The default security mechanism of Thunderbird is S/MIME. Go to www.mozilla-enigmail.org/forum/viewtopic.php?f=7&t=67 for more details.

GPG and e-mail

The first step is to create your own pair of keys for your e-mail account, user@example.com. It is simple. Just give the following command:

```
gpg --gen-key
```

You will be asked a few

questions and if in doubt, just use the defaults. It is better if you give a passphrase to protect your private key, especially if others may have access to the system you are using.

You will need to send your public key to your collaborators. So, export it as a text file and e-mail it:

```
gpg -a --export user@example.com >
my_public_key.asc
```

Your friends can call you to verify that the fingerprint of the key is valid. You can find out the fingerprint by the following command:

```
gpg --fingerprint
```

You can now sign and send an e-mail to your friends. In Evolution, choose the option 'Security' on the menu bar. Select the 'PGP Sign' check box.

When you get the public key from your friends and collaborators, you will need to import it. This step is also simple:

```
gpg --import his_public_key.asc
```

GPG expects each key to be signed by a trusted entity before it is regarded as valid. So, you will need to sign the key you have just imported as follows—assuming that your friend's e-mail address is friend@example.com:

```
gpg --sign-key friend@example.com
```

Now, you can encrypt the e-mail you are sending to your friend. When composing an e-mail, choose the 'Security' option from the menu bar and select the 'PGP Encrypt' check box. You can encrypt and sign the e-mail by selecting both the sign and the encrypt check boxes.

If you have received an encrypted and signed e-mail from your friend, Evolution will display it as usual, except that there will be a message at the bottom of the e-mail informing you that it had a valid signature and was encrypted.



**Hundreds of hardcore
Industry Professionals
Vouch for it.**



- Pioneer in **LINUX** training since 1999.
- 100% results every month.
- Many centurm students.
- 96 hrs. of training **PLUS** 36 hrs. of advance **LINUX** admin. training

**Next RHCE Exam on
09.03.09 & 23.03.09**

**Don't wait.
Register Now.**

Let the success vouch for you!

VECTRA
TECHNOSOFT **ADVANTAGE PRO**
IT TRAINING DIVISION OF
Vectra Technosoft Pvt. Ltd.
Networking with success

Wing 1& 2, IV Floor, Jhaver Plaza,
1A, Nungambakkam High Road,
Nungambakkam, Chennai - 600 034.
Phone : 28263529 / 30
E-mail : enquiry@vectratech.in
www.vectratech.in

If you forward the encrypted mail to someone else, including your friend, the recipient will not be able to decode the mail. This is very useful when sending e-mails to a relation who loves to gossip and has an uncontrollable mailing list! A side effect is that unless you copy an encrypted mail to yourself, you can't see what you sent.

If you do not have the public key of a recipient and you give the request to encrypt the mail, Evolution will give you an error. However, if both the Kirans – the union leader and the CFO in our opening paragraph -- had a public key in the key ring, encryption is not going to prevent you from making a mistake.

An example of an application

Programs can make mistakes—and they do so consistently. They do not normally make silly mistakes unless, of course, programmed to do so.

Suppose you want to send salary slips to all your employees, and want each employee to view only his or her salary details. Every employee, on joining, can create a key pair and register the individual public key with the company. The admin staff need not manage these keys securely! In fact, they can freely distribute the public key to anyone who needs it -- for example, the bank where a salary account is opened.

Python has a module called *pygpgme*, which is a wrapper for the *gpgme*, GPG Made Easy, library. It is installed on Fedora, as *Yum* needs it. It lacks one small thing—documentation!

The *gpgme* library is documented, but seems to lack any tutorials or articles on how to get started with it.

The solution in such cases is to download the source. You can actually ignore the source code and search for the test cases. That can act as an excellent starting point.

Encrypting/decrypting a file

For your application, you need to be able to encrypt a file. So, try the following code:

```
import gpgme
infile = open('salary_slip.txt')
outfile = open('salary_slip.asc', 'w')
ctx = gpgme.Context()
ctx.armor = True
recipient_key = ctx.get_key('friend@example.com')
ctx.encrypt_sign([recipient_key], gpgme.ENCRYPT_ALWAYS_TRUST, infile,
outfile)
outfile.close()
infile.close()
```

The code is pretty straightforward. Open the two files and obtain the GPG context. The 'armor' option creates an ASCII file rather than a binary one. Obtain the key by using the recipient's e-mail address, then

encrypt and sign the file by passing a list of the keys. The second option informs you that the keys should be trusted. You will be prompted for the passphrase while signing in, if you have specified one while creating your key.

The code for decrypting a file is even simpler:

```
import γπγμε
ινφιλε = οπεν(□σαλαρψ_σλιπ2.ασχ□)
ουτφιλε = οπεν(□σαλαρψ_σλιπ.ουτ□,□ω□)
χτζξ = γπγμε.Χοντεξτ()
σιγσ = χτζξ.δεχρψπτ_περιφψ(ινφιλε, ουτφιλε)
ουτφιλε.χλοσε()
ινφιλε.χλοσε()
```

gpgme will raise an exception in case decryption fails or the signature is not valid. The 'decrypt and verify' method will return a list of signatures. You may want to get some more information about the signatures.

Since there is only one key in your case, try the following code:

```
signing_key = ctx.get_key(sigs[0].fpr)
print signing_key.uids[0].name
print signing_key.uids[0].email
```

You get the key by using the fingerprint and then print the information you may need.

Let's suppose you just wanted to sign a text:

```
import gpgme
infile = open('salary_slip.txt')
outfile = open('salary_slip_signed.asc', 'w')
ctx = gpgme.Context()
ctx.armor = True
ctx.sign(infile, outfile, gpgme.SIG_MODE_CLEAR)
outfile.close()
infile.close()
```

You have chosen the clear sign mode so that the text is readable and the signature identifiable.

This will be enough for the moment. You can read the code in the tests subdirectory of the *pygpgme* source [pypi.python.org/packages/source/p/pygpgme/pygpgme-0.1.tar.gz] to learn more.

Mime and PGP

You are now in a position to combine encryption with the e-mail module so that the hard part is done by the application, and the user can access secure information very conveniently. The format for a Mime-encrypted message is described in www.ietf.org/rfc/rfc3156.txt.

Start with the various modules that need to be imported:

```
import smtplib
import pgpme
from email import encoders
from email.mime.base import MIMEBase
from email.mime.multipart import MIMEMultipart
from StringIO import StringIO
```

StringIO is a file-like class for manipulating a string buffer. It is, essentially, a memory file.

You will need to create a multi-part Mime formatted message with the attachment you wish to e-mail (see docs.python.org/library/email-examples.html for more details). Assume that you are attaching a PDF file:

```
def create_message(filename):
    outer = MIMEMultipart()
    fp = open(filename, 'rb')
    msg = MIMEBase('application', 'pdf')
    msg.set_payload(fp.read())
    fp.close()
    encoders.encode_base64(msg)
    msg.add_header('Content-Disposition', 'attachment', filename=filename)
    outer.attach(msg)
    return outer.as_string()
```

You will next encrypt the message:

```
def encrypt_payload(in_msg, out_msg):
    ctx = pgpme.Context()
    ctx.armor = True
    recipient_key = ctx.get_key('friend@example.com')
    ctx.encrypt([recipient_key], pgpme.ENCRYPT_ALWAYS_TRUST, in_msg,
out_msg)
```

Now, you will need to create another multi-part Mime message that has the encrypted content as the payload. The Mime body must consist of exactly two parts, the first with the content type “application/pgp-encrypted”. This part contains the control information. The second part contains the encrypted content as an octet-stream.

```
def mime_pgp_message(fp):
    outer = MIMEMultipart(_subtype='encrypted', protocol='application/
pgp-encrypted')
    outer['Subject'] = 'Attached Encrypted - 5'
    outer['To'] = 'friend@example.com'
    outer['From'] = 'user@example.com'
    msg = MIMEBase('application', 'pgp-encrypted')
    outer.attach(msg)
    enc_part = MIMEBase('application', 'octet-stream', name='encrypted.
asc')
    fp.seek(0)
    enc_part.set_payload(fp.read())
    outer.attach(enc_part)
    return outer.as_string()
```

Now, you are ready to send the message:

```
def send_message(sender, recipients, composed):
    s = smtplib.SMTP()
    s.connect()
    s.sendmail(sender, recipients, composed)
    s.close()

You would be calling the above routines as follows:
in_msg = StringIO(create_message('Open.pdf'))
out_msg = StringIO()
encrypt_payload(in_msg, out_msg)
composed = mime_pgp_message(out_msg)
send_message('user@example.com', ['friend@example.com'], composed)
```


Unfortunately, signing the document introduces one more level of complexity. Before encrypting the message, you would need to sign it. For this, too, a multi-part message with two parts in the body, is required. So, the steps would be:

1. Create the Mime message
2. PGP Sign the Mime message
3. Create a multi-part Mime message with protocol application/pgp-signature
4. PGP encrypt the signed Mime message
5. Create a multi-part Mime message with protocol application/pgp-encrypted
6. Send the message

Final words

This article turned out to be much harder to write than I expected, as I could not find any tutorials or simple documentation on using *pgpme* or Mime/PGP-encrypted, whether for Python or any other language. In case anyone knows of any, I would love to hear about it.

It is a pity that banks force us to change our passwords every few months. We also need to ensure that our passwords are not the same at all sites. Nor the same as the ones used on the previous few occasions. In short, keeping track of passwords is one horrendous problem. Desktop tools that help use the appropriate password for each application or website are a solution for this problem.

A public key environment would, unambiguously, shift the task of preventing any leakage of passwords from the host sites to the user only. The critical advantage is that we need to protect only one private password. Last but not least, it will save us from making hundreds of enemies because we used our Gmail password at a social networking site, involuntarily inflicting our friends with the “I want to be your friend” spam. **END** 

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com

Let a Thousand Languages Bloom!

Transifex can be your gateway to translations.

For most open source projects, translations are pretty important. Projects that are used by desktop users, such as desktop environments, GUI applications, and distributions, most frequently ship localised user interfaces, documentation, websites and other types of resources.

Take Fedora, for example -- one of the most popular Linux distributions out there. Around 60 per cent of its users use a localised desktop, and the percentages may probably be higher with other major desktop environments. In the case of Fedora content, this gets translated to something like 3-5 million users. With contributions having such a large audience

and impact, it's no surprise that the open source translation community is very active, and most major open source projects enjoy an active community devoted to translating the project into various languages.

Challenges in FOSS localisation

Typically, software developers use an internationalisation platform like *gettext*, which parses the source code and extracts the translatable strings from the code into special PO files. These files are handed to translators, who translate them into a target language using a variety of tools.

The challenge for most projects lies in receiving those translation files back in

their version control system (VCS). Giving access to your VCS to a few developers is usually okay, but having to administrate accounts for hundreds of translators could be a challenge. To avoid that, some developers even decide to only accept translations with bug reports or e-mail attachments. But a developing product usually means that “strings are changing often”, and with each release, translators will send a new batch of translations in. That’s a lot of bug reports and e-mails.

Larger projects usually have the advantage of developing their own translation community. In which case, however, some developers feel more productive using a different type of VCS, and some others even host their project on external servers. The consequences of these approaches are either low productivity, or just a small number of translators and quality that suffers.

Finding a solution

Transifex has been developed as a solution to these issues, and to make translations dead-simple both for developers and translators. The goal with Transifex was to work as a translation proxy and handle the mechanical processes for both these groups of users, allowing them to work more efficiently and effectively.

Developers give Transifex access to their source repository. The Transifex “robot” can log in to a number of different versioning systems and grab the translation files for the translators. The latter log in to a unified, easy-to-use interface, independent of the upstream VCS type and location, and receive the translations they need. Upon translation, they can use the same interface to submit the files back to the VCS.

How it works

Richard Hughes is the software developer of PackageKit. He hosts his project in *packagekit.org*, and

needs to find a way to receive quality translations in a hassle-free way. He fires up his browser to an existing Transifex server (such as the soon-to-be-launched *transifex.net*) and registers his project there. He then receives an SSH key and uses it to create a special user on his server, with write access in the translation directories. His project is now ready to receive translations.

At this point, Richard is asked whether he’d like Transifex to scan its translation memory from other projects to bootstrap the translations of his own projects. He’s delighted to see that his PO files have been translated to somewhere between 20-40 per cent with no human interaction.

Piotr is a Polish translator who loves translating free software GUIs. He has registered with Transifex and requested to receive notifications for new projects registered, which might interest him. He receives an e-mail with a direct link to the Polish PackageKit translation and another link that he can use to submit the file back.

Once the file is submitted back, Richard is notified that language translation for Polish is now at 100 per cent.

Architecture details

Under the hood, Transifex abstracts all VCSs and runs a clone/checkout on the repository. It identifies the i18n method and the translation files. Depending on the i18n method, it compares the translation files with the template file (for example, the English one) and calculates translation statistics for each one.

The management burden is removed from developers, who can concentrate on what they do best, which is writing code. Translators can use their single Transifex login account to contribute to any project they like, as long as it’s registered on Transifex.

As a high-level Python application, the service includes hooks that can improve the

NACE™
move to the next level

TRAINING & CONSULTANCY

**BECOME AN
OPEN SOURCE
PROFESSIONAL.**

**IT INDUSTRY IS
WAITING FOR...**



**LINUX
APACHE
MySQL
PHP**

Linux Basics

Linux Administration

Linux Networking

Linux Security

OpenOffice.org 3.0



CHENNAI

T.NAGAR : 99405 42222

ANNA NAGAR : 97909 25554

ADYAR : 97909 25550

MYLAPORE : 90030 33233

TAMBARAM : 04442059008

AMBATTUR : 94453 22869

PERAMBUR : 93808 76090

VILLIVAKKAM : 9940233360

VELLORE : 9994241383

SALEM : 9944175120

PONDICHERRY : 97918 53774

BANGALORE : +91 9940022277

HYDERABAD

MALAKPET : 04032442998 / 99

ASHOK NAGAR : 9703333839

PANCHAGUDA : 040 64562045 / 46

www.naceedu.com

workflow in a number of ways. Pre-commit, the validity of the file's syntax is checked, avoiding breaking the developer's build process with broken files. It also allows fine-grained permissions to files the translators need access to. Post-commit, Transifex can notify language leaders and others about file submissions, provide RSS feeds for submissions, etc.

Transifex currently supports git, hg, cvs, svn and bazaar, and adding more VCSs is a matter of writing a few lines of code. Its developers serve POT-based projects, and are looking forward to extending the i18n support to include intltool-based projects (GNOME), XLIFF, etc. The login mechanism also supports OpenID.

Development of Transifex

The development of Transifex began as part of the 2007 Google Summer of Code project by myself (Oh! Hi! I'm Dimitris Glezos :-)). It was initially written in Python using the TurboGears framework, and right after the summer it was put into production in Fedora, used by more than 100 projects and 500 translators.

Next year, Transifex was presented in more than 10 international conferences, including FOSS.in 2008. In the summer, Transifex earned three more GSoC applications and was re-written from scratch using the Django Python framework, now including many of the suggestions from existing users. Development has taken place since then on transifex.org and on the [transifex-devel](mailto:transifex-devel@lists.fedoraproject.org) mailing list.

In the meantime, other projects liked the platform and joined in our efforts. GNOME's Damned Lies and Vertimus tools migrated their code to Django, with the goal of being merged with Transifex at some point in the future.

Future features

With more contributors joining in the developer team, Transifex is now moving towards a stabilised platform to serve independent and upstream software projects and then on to bigger ones.

One of the immediate features we'd like to add is per-VCS file monitoring, so that translators can 'track' a project and get notified when the translation percentage for their language changes. Adding commenting support for projects and submissions, as well as developing support for file uploads will enable translators to better collaborate in QA.

Another often requested feature is the development of a command-line interface allowing translators to do something like the following:

```
$ tx set-language bn_IN
$ tx get-collection Fedora
Received anaconda/po/bn_IN.po
Received packagekit/po/bn_IN.po
$ # Translation...
```

```
$ tx send-collection Fedora
Sent 'anaconda/po/bn_IN.po' (100% translated)
```

The vision: Transifex.net

As mentioned earlier, Transifex allows downstream communities to send files directly to the VCS of upstream projects. One might wonder then, which Transifex community should an independent project choose to receive translations from—Fedora, GNOME, or example.com?


Having a common place where open source translations take place is key to link translation communities together and reach new levels of collaboration between translation teams. Here's a plan we're evolving with www.transifex.net: Establish a healthy network where developers can translate their applications and translators can contribute to their favourite projects. Project teams that wouldn't like to undertake the trouble of setting up their own Transifex instance, should have a stable, rich-in-features service, to join their efforts with the rest of the open source community, under a common umbrella.

Becoming a contributor

Transifex is written in Python and utilises the awesome Django branch with its infamous top-notch documentation. This makes it really easy for folks to join in and extend the platform with the features they'd like to see added. Development information can be found at transifex.org/wiki/Development. To set up a development environment of your own, check out the documentation at docs.transifex.org/intro/install.html.

An example of an easy task would be to add support for associating registered projects with their maintainers/developers. This will give translators a contact point for more information on the project and for conflict resolution. Creating a patch that adds simple support for project maintainers is a matter of a few lines of code: add a foreign key from the *Project* model to the *User* and probably edit the *User Profile* page to include a section listing the projects the user maintains.

Adding support for more VCSs and i18n back-ends is also quite feasible because of the abstractions Transifex includes in those areas. For most needs, one just copies a Python file and changes accordingly. We've marked quite a few tickets with the 'easy_task' keyword, so check out transifex.org/report/9 to start hacking.

Let a thousand languages bloom!  **END**

By: Dimitris Glezos

The author is a member of the Fedora Board and the current Fedora Localisation Leader. He loves watching people read content in their native language, and so he founded Indifex, a fresh start-up specialising in quality and effective translation services. He lives in Greece and likes hacking, photography and rock climbing.

An Exclusive B2B Exhibition - With a Senior Level Conference

Connect **2009**

**4th Information & Communications
Technology Exhibition & Conference**

5 - 7 May 2009

Karachi Expo Centre

www.connectit.com.pk



ORGANISED BY



Pegasus Consultancy (Pvt) Ltd

2nd Floor, Business Centre,
Mumtaz Hassan Road, Karachi - Pakistan
Tel : +(92 21) 111 734 288 (PEGCON)
Fax : +(92 21) 241 0723

E-mail : info@connectit.com.pk, URL : www.pegasus.com.pk

ACTIVITY PARTNER



National Response Centre
for Cyber Crimes (NRC3)

PARTNER SHOW



Telecom & Broadcasting

SUPPORTED BY





KVM

Virtualisation, the Linux Way

KVM, the Kernel Virtual Machine monitor, was announced in late 2006, and was merged in Linus' tree in December the same year. It has very quickly gained wide acceptance and adoption for being the most promising and capable virtualisation strategy on Linux. Though a very young project, new features are being added at a very brisk pace thanks to the interest taken by several companies and developers across the globe.

*B*efore we look at KVM, let's go over what virtualisation is, in the context of the technologies available and then move on to what makes KVM different and how easy it is to use.

'Virtualisation' means the simulation of a computer system, in software. The virtualisation software creates an

environment for a 'guest', which is a complete OS, to execute within this created world. This means that the view that should get exported to the guest should be of a complete computer system, with the processor, system peripherals, devices, buses, memory and so on. The virtualisation software can be strict about what view to export to the guest -- for

example, the processor and processor features, types of devices, buses exported to the guest -- or it can be flexible, with the user getting a choice of selecting individual components and parameters.

There are some constraints to creating a virtualised environment. A set of sufficient requirements noted by Popek and Goldberg in their paper on virtual machine monitors are:

- **Fidelity:** Software that is running in a virtualised environment should not be able to detect that it is actually being run on a virtualised system.
- **Containment:** Activities within a virtual machine (VM) should be contained within the VM itself without disturbing the host system. A guest should not cause the host, or other guests running on the host, to malfunction.
- **Performance:** Performance is crucial to how the user sees the utility of the virtualising environment. In this age of extremely fast and affordable general-purpose computer systems, if it takes a few seconds for some input action to get registered in a guest, no one will be interested in using the virtual machine at all.
- **Stability:** The virtualisation software itself should be stable enough to handle the guest OS and any quirks it may exhibit.

There are several reasons why one would want virtualisation. For data centres, it makes sense to run multiple servers (Web, mail, etc) on a single machine. These servers are mostly under-utilised, so clubbing them on one machine with a VM for each of the existing machines, makes way for fewer machines, less rack-space and lower electricity consumption.

For enterprises, serving users' desktops on a VM simplifies management, IT servicing, security considerations and costs by reduced expenditure on desktops.

For developers, testing code written for different architectures or target systems becomes easier, since access to the actual system becomes optional. For example, a new mobile phone platform can be virtualised on a developer machine rather than actually deploying the software on the phone hardware each time, allowing for the software to be developed along with the hardware. The virtualised environment can also be used as validation for the hardware platform itself before going into production to avoid costs arising later due to changes that might be needed in the hardware.

There are several such examples that can be cited for any kind of application or use. It's not impossible to imagine a virtualised system being beneficial anywhere a computer is being used.

Now is a good time to get acquainted with some terms (the mandatory alphabet soup) that we'll be using throughout the article:

- **VM:** Virtual Machine
- **VMM:** Virtual Machine Monitor
- **Guest OS:** The OS that is run within a virtual machine
- **Host OS:** The OS that runs on the computer system
- **Paravirtualised guest:** The guest OS that is modified to have knowledge of a VMM
- **Full virtualisation:** The guest OS is run unmodified in this environment
- **Hypervisor:** An analogous term for a VMM
- **Hypercall:** The medium that a paravirtualised guest and the VMM communicate on

Types of VMM

There are several virtual machine monitors available. They differ in various aspects like scope, motivation, and method of implementation. A few types of monitor software are:

- **'Native' hypervisors:** These VMMs have an OS associated with them. A complete software-based implementation will need a scheduler, a memory management subsystem and an IO device model to be exported to the guest OS. Examples are VMWare ESX server, Xen, KVM, and IBM mainframes. In IBM mainframes, the VMM is an inherent part of the architecture.
- **Containers:** In this type of virtualisation, the guest OS and the host OS share the same kernel. Different namespaces are allocated for different guests. For example, the process identifiers, file descriptors, etc, are 'virtualised', in the sense a PID obtained for a process in the guest OS will only be valid within that guest. The guest can have a different userland (for example, a different distribution) from the host. Examples are OpenVZ, FreeVPS, and Linux-Vserver.
- **Emulation:** Each instruction in the guest is emulated. It is possible to run code compiled for different architectures on a computer, like running ARM code on a PowerPC machine. Examples are qemu, pearpc, etc. qemu supports multiple CPU types, and runs ARM code under x86 as well as x86 under x86, whereas pearpc only emulates the PPC platform.

Hardware support on x86

Virtualising the x86 architecture is difficult since the instruction and register sets are not compatible with virtualisation. Not all access to privileged instructions or registers raises a trap. So we either have to emulate the guest entirely, or patch it at run-time to behave in a particular way.

With the two leading x86 processor manufacturers, Intel and AMD, adding virtualisation extensions to their processors, virtualising the x86 platform seamlessly has become easier. The ideas in their virtualisation extensions are more or less the same, with the implementation, instructions and register sets being

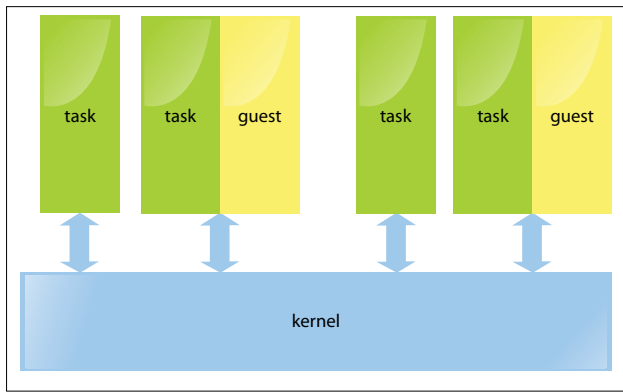


Figure 1: The KVM process model

slightly different.

The new extensions add a new mode, 'guest mode', in addition to the user-mode and kernel-mode that we had (ring -1 in addition to the rings 0-3, with the hypervisor residing in ring -1). The implementations also enable support for hiding the privileged state. Disabling interrupts while in guest mode will not affect the host-side interrupts in any way.

The KVM way

For KVM, the focus is on virtualisation. Scheduling of processes, memory management, IO management are all pieces of the puzzle that were already available. It didn't make sense writing all those components again, especially since the software available was in wide use, was very well tested (hence stable), and written by experts in the field.

The way KVM solves the virtualisation problem is by turning the Linux kernel into a hypervisor. Scheduling of processes and memory management is left to Linux. Handling IO is left to qemu, which can run guests in userspace and has a good device model. A small Linux kernel module has been written and it introduces the 'guest' mode, sets up page tables for the guest and emulates some instructions.

This fits nicely into the UNIX mindset of doing one thing and doing it right. So the *kvm* module is all about enabling the guest mode and handling virtualised access to registers. From a user's perspective, there's almost no difference in running a VM with KVM disabled and running a VM with KVM enabled, except, of course, the significant speed difference (Figure 1).

The philosophy also extends to the development and release philosophy Linux is built on: release early and often. This allows for fast-paced development and stabilisation. Developers can track the latest and greatest codebase and keep enhancing it.

The latest stable release is part of Linux 2.6.x, with bug fixes going in 2.6.x.y. The KVM source code is maintained in a git tree. To get the latest KVM release or the latest git tree, head to kvm.qumranet.com for download details.

Using KVM

First, since KVM only exploits the recent hardware advances, you should make sure you have a processor that supports the virtualisation extensions. This can be quickly found out by using the following command:

```
egrep '^flags.*(vmx|svm)' /proc/cpuinfo
```

If there's any output, it means the necessary capability to run KVM exists on the CPU. If you have a CPU that's less than three years old, you're set.

You now need to run a recent 2.6 Linux kernel.

If you already run a recent Linux kernel with KVM either compiled in the kernel or compiled as modules, you can work with it if you don't want to compile the modules yourself. All the distributions these days ship *kvm* modules by default, so getting KVM is just a matter of fetching the necessary packages from your distribution's site if they're not already installed.

There are two parts to KVM: the kernel modules and the userspace support, which is a slightly modified version of qemu.

You can also download the KVM sources from *kvm.qumranet.com/kvmwiki/Downloads*.

Building the userspace utilities from the tarball needs to have a few libraries present. The detailed list and instructions are given at *kvm.qumranet.com/kvmwiki/HOWTO*.

Once you have the kernel module and the userspace tools installed (either by building or installing from your distribution packages), the first thing to do is create a file that will hold the guest OS. We have previously seen how to do that using the virt-manager GUI in January. If you like using the command prompt, which is my preferred way of doing it, here's a quick tutorial:

```
$ qemu-img create -f qcow2 debian-lenny.img 10G
```

This will create a file called *debian-lenny.img* of size 10G in the 'qcow2' format. There are a few other file formats supported, each with their advantages and disadvantages. The qemu documentation has details.

Once an image file is created, we're ready to install a guest OS within it. First, insert the *kvm* kernel modules in the kernel if they have been compiled as modules.

```
$ sudo modprobe kvm
```

```
$ sudo modprobe kvm-intel
```

...or:

```
$ sudo modprobe kvm-amd
```

On Debian-based systems you can add yourself to

the 'kvm' group:

```
$ sudo adduser user kvm
```

Logout and login again, and you can then start the VMs without root privileges.

```
$ qemu-system-x86_64 -boot d -cdrom /images/debian-lenny.iso \
-hda debian-lenny.img
```

On Fedora, this is:

```
# qemu-kvm -boot d -cdrom /images/debian-lenny.iso \
-hda debian-lenny.img
```

This command starts a VM session. Once the install is completed, you can run the guest with the following command:

```
$ qemu-system-x86_64 debian-lenny.img
```

You can also pass the *-m* parameter to set the amount of RAM the VM gets. The default value is 128 M.

Recent KVM releases have support for swapping guest memory, so the RAM allocated to the guest isn't pinned down on the host.

Troubleshooting

There will be times when you run into some bugs in VMs with KVM. In such cases, there will be some output in the host kernel logs. That can help you search for similar problems reported earlier and any solutions that might be available. In most cases, upgrading to the latest KVM release and running it might fix the problem.

In case you don't find a solution, running the VM by passing the *-no-kvm* command line to qemu will start it without KVM support. If this also doesn't solve the problem, it means the problem lies in qemu itself.

Another thing to try out might be to pass the *-no-kvm-irqchip* parameter while starting a VM.

You can also ask on the friendly *#kvm* IRC channel at Freenode, or on the *kvm@vger.kernel.org* mailing list, and someone will help you out.

qemu monitor

The qemu monitor can be entered by typing the *Ctrl+Alt+2* key combination when the qemu window is selected, or by passing *-monitor stdio* to the qemu command line. The monitor gives access to some debugging commands and those that can help inspect the state of the VM.

For example, info registers show the contents of the registers of the virtual CPU. You can also attach USB devices to a VM, change the CD image and do other interesting things via the qemu monitor.

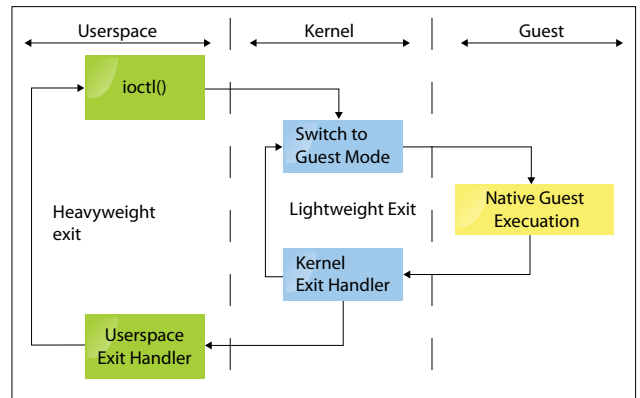


Figure 2: The KVM execution model

Migration of VMs

Migrating VMs is a very important feature for load-balancing, hardware upgrades, and software upgrades for zero or very minimal downtime. The guests are migrated from one physical machine to another, and the original machine can then be taken down for maintenance.

The advantage in the KVM approach is that the guests are not involved at all in the migration. Also, nothing special needs to be done to tunnel a migration through an SSH session, to compress the image being migrated, etc. You can even pass the image through any program you want before it's transmitted to the target machine. The UNIX philosophy holds good here as well. Also, unless specific hardware or host-specific features are enabled, the migration can be made between any two machines. Moreover, stopped guests can be migrated as well as live guests. We add the migration facility within qemu, so no kernel-side changes are needed to enable it. The device state-syncs to achieve migration, and the VM state is seamlessly provided and managed within userspace.

On the target machine, run qemu with the same command-line options as was given to the VM on the source machine, with additional parameters for migration-specific commands:

```
$ qemu-system-x86_64 -incoming <protocol://params>
```

Example:

```
$ qemu-system-x86_64 -m 512 -hda /images/f10.img -incoming params
```

On the source machine, start migration using the 'migrate' qemu monitor command:

```
(qemu) migrate <migration-protocol://params>
```

An example of the source qemu monitor command:

```
(qemu) migrate tcp://dst-ip:dst-port
```

...while the command-line parameter of the target qemu migration is:

`-incoming tcp://0:port`

Migration support in qemu is currently being overhauled to a newer, simpler version, and more functionality is being added. Support for some previously-supported functionality, like migration via SSH or via a file isn't added yet in the new infrastructure—could be an interesting area for people looking to contribute to explore.

Advantages of the KVM approach

There are several advantages in doing things the Linux way: we reuse all the existing software and infrastructure available, and no new commands need to be learned and not many need to be introduced. For example, `kill(1)` and `top(1)` work as expected on the guest task on the host system. The guests are scheduled as regular processes. As we saw in Figure 1, each guest consists of two parts: the userspace part (qemu) and the guest part (the guest itself). The guest physical memory is mapped in the task's virtual memory space, so guests can be swapped as well. Virtual processors in a VM are merely threads in the host process.

When KVM was initially written, the design parameters were to support x86 hosts, focus only on full virtualisation (no modifications to guest OS) and with no modifications to the host kernel. However, as KVM started gaining developers and interesting use-cases, things changed. Because of the simple and elegant solution, developers took a liking to the approach and new architecture ports for s390, PowerPC and IA-64 were added within months of starting the ports. Paravirtualisation support also has been added, and pv drivers for net and block devices are available. If a guest OS can communicate with

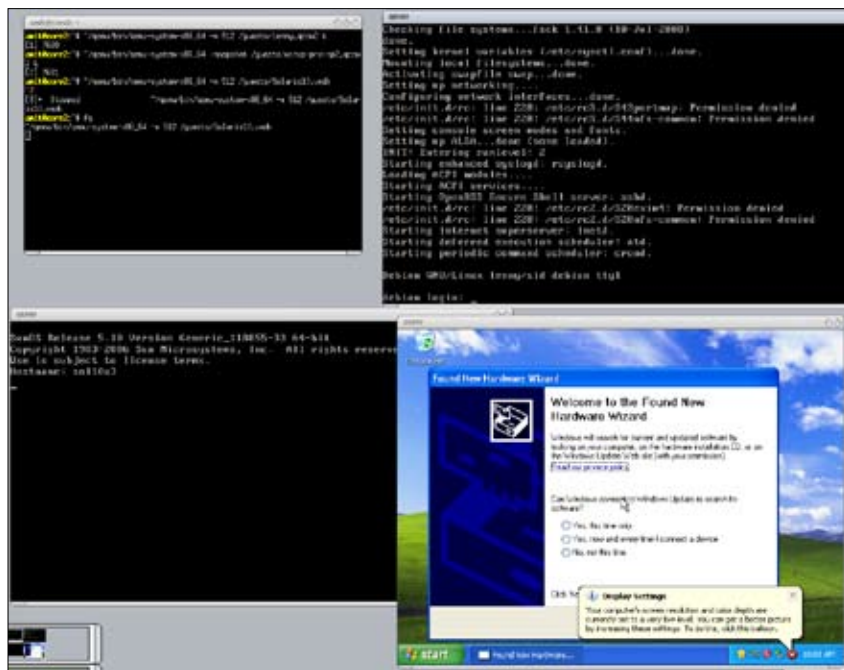


Figure 3: Running multiple operating systems using KVM


the host, several activities can be speeded up, like network activity or disk IO. Also, modifications to the host OS (Linux) that improve scheduling and swapping, among other things, were proposed and accepted.

KVM seamlessly works across all machine types: servers, desktops, laptops, and even embedded boards. Let's look at each, one by one.

- **Servers:** There's the distinct advantage of being able to use the same management tools and infrastructure as Linux uses. KVM integrates with the Linux scheduler, IO stack, all available filesystems and supports live migration; it has ready support for NUMA and 4096-processor machines.
- **Desktops/laptops:** KVM works on anything Linux works on. The normal desktop doesn't change. You can suspend/resume work as expected, even while virtual machines are running.
- **Embedded:** Linux already supports lots of boards, architectures and machine types. Real-time scheduling

is supported. And if you're wondering why one would want to use virtualisation on an embedded machine, here are a few of the many reasons: to sandbox untrusted code, reliable remote kernel upgrades, uniprocessor software on multiprocessor cores, running legacy applications, etc.

All this highlights the differences between KVM and some of the other virtualising solutions.

There's a lot of ongoing work to stabilise KVM and improve the performance of the block and net paravirtualised drivers. There is also some development activity in the regression test suite framework that's ongoing. If you're looking to contribute to KVM, these areas are waiting for you! **END** 

By: Amit Shah

The author has been working on systems software on the Linux kernel for seven years now. He considers himself to be fortunate enough to be accepted as the first off-site developer for the KVM team when he joined Qumranet in 2007. He's worked on a few interesting problems in KVM and is now part of the bigger virtualisation group in the Red Hat family.

Building A Highly Available Nginx Reverse-Proxy Using



Heartbeat

Last month we discussed how to set up a highly available cluster of Web servers that are load balanced using nginx. One shortcoming in that set-up was the reverse-proxy server itself, which on crashing, will cause the entire Web server cluster to go down. Therefore, we would need to build high-availability in the reverse-proxy server itself.

A cluster in computing is a term used to describe a group of closely linked computers often appearing as a single entity to the outside world. There are various types of clusters -- high-availability (HA) clusters, load-balancing clusters, compute clusters or high performance computing (HPC) clusters and grids.

An HA cluster is also known as a failover cluster and is typically meant to improve service availability rather than performance, by using redundant nodes. There are many models of HA cluster configuration such as active-passive, active-active, N+1, N+M, N-to-1 and N-to-N.

Load-balancing clusters distribute the workload evenly among various redundant nodes. There are various different algorithms, using which the load-balancer distributes the load among member servers.

HPC clusters are used for highly CPU-intensive compute jobs. There are various types of compute clusters. The common distinguishing factor is the coupling of the compute nodes. Typically, there are specialised scientific applications that run on these clusters. The applications are built

using libraries supporting parallel processing. A popular example of a compute cluster is the Beowulf Clusters. These are built using commodity hardware and run on FOSS systems like FreeBSD or GNU/Linux. Typically, a Beowulf cluster uses either MPI (Message Passing Interface) or PVM (Parallel Virtual Machine) libraries that allow a programmer to divide a task among the nodes of the cluster, and then recollect and assemble the results later on.

A grid is a special class of compute clusters with possibly heterogeneous nodes that are not so tightly coupled with each other. All nodes in the grid target single problems that require a great number of CPU cycles and a large amount of data. A grid typically divides the entire computation work into jobs that are independent and do not share data with each other. The intermediate results of one job in the grid do not affect the other jobs running on other nodes.

Heartbeat is a piece of software from 'The High Availability Linux' project, which provides high-availability clustering solutions for a wide range of *nix operating systems, including (though not limited to) GNU/Linux, FreeBSD and OpenBSD.

Primary/secondary node IP addresses		
Parameters / Node	rproxy1.unix-clinic.net	rproxy2.unix-clinic.net
eth0 (192.168.1.x is Private Subnet for Heartbeat)	192.168.1.1	192.168.1.2
Eth1 (administrative address)	172.202.2.1	172.202.2.2
Service address 10.8.0.1	Primary node	Secondary node

Table 1

The architecture: active-passive HA cluster

The two primary modes of an HA cluster are:

- **Active-passive:** In Active-Passive HA clusters, the primary node is active and serves the request. In case it fails, then the services are transferred to the secondary node, either through automatic or manual failover.
- **Active-active:** In active-active HA clusters, both the nodes remain active all the time and serve their respective requests. In case one of the nodes goes down, the services running on that node are failed-over to the other node in the cluster. This way an active-active HA cluster is used when you have multiple services under the high-availability requirement.

The service being served by the HA cluster depends on the IP address, so we first need to distinguish between the 'administrative address' and the 'service address'. Each interface on the cluster nodes should have an administrative address and can optionally have one or more service addresses, depending on the cluster configuration (active-active or active-passive) and state (active or standby).

An administrative address is one that is in control of the operating system and is brought up and down with the OS. A service address is one that is under the control of the Heartbeat software, which then controls its allocation to one of the cluster nodes. The node where this service address should reside, by default, is known as the primary or the active node, and the other node in the cluster is known as a secondary, failover or passive node.

In failover clustering, when a failover happens, the secondary node takes over the service address and becomes active. This is how we will be configuring our cluster.

In our case, the service being offered by the cluster, a reverse-proxy server, depends on the IP address. So we need to take care of the following points:

- Make sure that the nginx server is not started automatically on any node, but is under the control of Heartbeat.
- The functioning of nginx will be dependent on the availability of the service IP address and hence, in case of a failover, we need to make sure that nginx is started after the service IP address has been taken over by the secondary node.

Table 1 can be referred to for setting up the networking. For a test environment, a flat network would do.

Installation and configuration of Heartbeat

On Debian-based systems, Heartbeat can be installed as follows:

```
# apt-get install heartbeat-2
```

On CentOS 5.2, after subscribing to the 'extras' repository, execute the following command:

```
# yum install heartbeat
```

Typically, in an active-passive HA environment, the nodes of the cluster have an identical set-up, so unless otherwise noted, all the configurations have to be done identically on both the nodes. It is not mandatory that the nodes are of identical hardware configuration, but this is recommended for a production environment. Also, having the same OS on both the nodes will help from the maintenance and troubleshooting point of view.

In the high-availability environment, all nodes should be able to see one another, irrespective of the availability of the DNS; so we need to make sure that there are relevant entries in the `/etc/hosts` file on each node. My host's file for this set-up looks like what's shown below:

```
# cat /etc/hosts
.....
192.168.1.1 rproxy1
192.168.1.2 rproxy2
```

Now you should be able to *ping* both nodes from each other. This will make sure that the Heartbeats from both nodes will see each other, irrespective of the DNS. You can build as many redundancies as you want; e.g., you can have bonding for the interface that's responsible for sending Heartbeat, you can have trunking on the switch ports, and so on. Whatever you do, just make sure it's not an overkill and the set-up is not unnecessarily expensive.

The *ha.cf* file

The main configuration file for Heartbeat is *ha.cf*, which lists the nodes of the cluster, communications topology and all the features that are enabled. The order of directives in *ha.cf* matters, so make sure you take note of that.

The minimum *ha.cf* file will look like the following, on both the nodes:

```
# cat /etc/ha.d/ha.cf
use_logd on
keepalive 2
deadtime 30
warntime 10
initdead 120
udpport 695
ucast eth0 192.168.1.1
ucast eth0 192.168.1.2
# bcast eth0
```

auto_failback on
node rproxy1 rproxy2

The explanation of these options is given below:

- **use_logd** specifies the use of the logging daemon to log all the messages. This option has deprecated the debugfile/logfile/logfacility log options. Using this is recommended.
- **keepalive** specifies the number of seconds between two heartbeats.
- **deadtime** specifies the number of seconds after which a host is considered dead if not responding.
- **warntime** specifies the number of seconds after which the late Heartbeat warning will be issued.
- **initdead** is the number of seconds to wait for the other host after starting Heartbeat, before it is considered dead.
- **udpport** is the port number used for bcast/ucast communication. The default value for this is 694, but I've used 695 because I had a pair of HA LM-1500 load balancing appliances from Kemp Technologies on the same subnet for some Citrix servers. This appliance is based on Linux and to me it looked like it was using Heartbeat for HA on the default UDP port 694.
- **bcast/ucast** is the interface on which to broadcast/unicast. If you are planning to make this only a two-node cluster, then there is no need for sending broadcasts; instead, use unicast. Now you will see that one of the IP addresses to which unicast is being sent is the local machine itself. I have done this to make sure that the *ha.cf* file is identical on both the cluster nodes. The unicast directives sent to the local machines are effectively ignored.



Note: If you have changed the default UDP port (like I have done above) then make sure that the 'ucast' or 'bcast' line is after the 'udpport' line; otherwise, the default port 694 will be used. Remember that the order of the DIRECTIVES in *ha.cf* is important and matters.

- If **auto_failback** is set to 'on', then resources are automatically failed back to its primary node.
- The **node** directive specifies the nodes in the HA set-up. The name specified here must match with the **uname -n** of the cluster node.

The *haresources* file

Now we need to tell Heartbeat about the resources the cluster will be managing. There are two ways of doing this. One, by using the *haresources* file, and two, by enabling the Cluster Resource Manager (CRM) and using *cib.xml*. In Heartbeat 2.x versions, if CRM is enabled, *haresources* is not used.

Setting up clustering using CRM is unnecessarily complicated for our simple set-up, although the Linux-HA project provides the command line tools and GUI tools to manage the set-up. In our set-up, we will use the Heartbeat R1 style clustering (named so because of its compatibility with the older release 1.x of Heartbeat). Once you can get this working, you can move towards setting up the

Heartbeat R2 style clustering.

If you are using the *haresources* file for your set-up, you need to make sure this file is identical on both the machines. The general syntax of this file is to list the preferred-node followed by the list of resources that will be running on this preferred-node. All the resources specified on a single line are called a resource-group. In order to continue on the next line, a backslash can be used. The first resource in each resource group (in case you are specifying multiple resource groups) needs to be unique because it is used as the resource-group name. The preferred-node is the one where the listed resources will run by default when both (or all) nodes of a cluster are available and the 'autofailback' option is set to 'on' in the *ha.cf* file.

Shown below is my *haresources* file:

```
# cat /etc/ha.d/haresources
rproxy1 10.8.0.1 nginx
```

While taking-over or acquiring the resources, Heartbeat uses left-to-right ordering and while releasing, a right-to-left order. All resources under cluster control must have a resource control script that is typically located in either */etc/init.d* or */etc/ha.d/resource.d* directories. Any script that can take at least two parameters, start and stop—to start and stop the resource, respectively—can be used as a resource control script. The IP address used in our *haresources* file is the service IP address, where our nginx reverse proxy server will be serving the requests. In DNS, this IP is to be mapped with *www.unixclinic.net*, which is what the customers will be accessing. You can see that we have not used any resource control script to acquire the IP address. The reason for this is that a service IP address is typically the requirement for all kinds of HA clusters that we set up and hence Heartbeat, by default, uses a resource control script called *IPaddr* for acquiring the IP address. So instead, the *haresources* file could have been written as:

```
rproxy1 IPaddr::10.8.0.1 nginx
```

...or if we want to specify netmask, interface or broadcast values for the service IP, then we can use the following syntax:

```
rproxy1 IPaddr::10.8.0.1/255.255.255.0/eth1/255.255.255.255
```

In our *haresources* file, we have only specified the service IP address for the cluster, but have not specified which interface in the machine will acquire this IP address; neither have we specified the netmask and broadcast values. In such cases, these values are set automatically by Heartbeat, by looking at the routing table. Heartbeat attempts to find the lowest cost route to the service IP address, and if multiple interfaces are found to provide the lowest cost route, then the first such route is considered. This basically means that the default route of the system is the least preferred. For setting up the broadcast address, the largest available address is used.

For details on this, see the side-box: 'IPaddr versus IPaddr2

Cluster Resource Manager'.

The second resource `nginx` is the name of the start-up script to start/stop the resource `nginx` and by default, Heartbeat looks for it in either `/etc/init.d` or `/etc/ha.d/resource.d`. The lines in `haresources` are translated as follows:

- On Heartbeat startup, acquire IP address 10.8.0.1 and then start `nginx` on node `rproxy1`
- When Heartbeat is stopped, stop `nginx` and then release IP 10.8.0.1 on node `rproxy1`

The *authkeys* file

The *authkeys* file is very important in maintaining the security of the cluster, as it authenticates the cluster nodes. This file should be owned by the root with permissions set to 600 (that is, readable and writeable by only the root), otherwise Heartbeat will refuse to start. Also, all the nodes of the cluster should have an identical *authkeys* file. Heartbeat supports three authentication methods: `crc`, `md5` and `sha1`. You probably don't use a serial or crossover connection for Heartbeat—if you do, `crc` is a good choice. For set-ups where the Heartbeat travels over the network, `sha1` is a good choice. We will use `sha1` in our set-up.

The following is the *authkeys* file in our set-up:

```
# cat /etc/ha.d/authkeys
auth 1
1 sha1 ThisIsMySecretKeyAndICanChooseAnyStringHere
```

Typically, the number of the key used is 1, but you can use any number ranging from 1-15. Just make sure that whatever number you use with the `auth` line, is present in one of the keys listed in the following lines. In order to generate a complete random number as a secret key for `sha1`, use the following command as suggested in the Linux-HA website and other places on the Web, and replace the string “ThisIsMySecretKeyAndICanChooseAnyStringHere” with the output of the following command:

```
# dd if=/dev/urandom count=4 2>/dev/null | openssl dgst -sha1
```

Controlling the start-up of cluster resources

Since our proxy server functionality is dependent on the IP address, we need to make sure that the `nginx` server does not get started automatically by the system start-up scripts during system boot. On Debian and its derivatives, this can be done as follows:

```
# invoke-rc.d nginx stop
# update-rc.d -f nginx remove && update-rc.d nginx stop 45 .
```

Note how the *update-rc.d* script is used above. If we would have just removed the `nginx` from start-up as done by the *remove* command, an update to `nginx` through *apt-get* would have triggered the creation of the missing symbolic links in the *rc?.d* directories. This is not what we wanted and

IPaddr versus IPaddr2 Cluster Resource Manager

Heartbeat uses either the `IPaddr` or `IPaddr2` resource manager script to configure IPv4 service addresses.

By default, the `IPaddr` script is used. The cluster resource manager scripts are located in the `/etc/ha.d/resource.d` directory. The basic syntax for both the scripts are:

```
IPaddr::ip-address[/netmask[/interface[/broadcast]]
IPaddr2::ip-address[/netmask[/interface[/broadcast]]
```

The difference between the two is that the `IPaddr` script uses the ancient method of IP aliasing, whereas `IPaddr2` uses the new IPv2 method of setting secondary IP address. There is a limit of 100 aliases on the `IPaddr` script, while there is no such limit on `IPaddr2` script. In case you use `IPaddr2` (which is preferable, and the default cluster resource manager script in Heartbeat now), you will not be able to see the acquired service address using the `ifconfig` command. This can be done by using the `ip addr show` command. For details on the `ip` command, refer to the “Linux Advanced Routing and Traffic Control” how-to link in the References section.

will bring our cluster down. For post installation/update scripts not to create or update symbolic links, and in order to leave `nginx` in the default disabled state, we have created a stop symlink in run-level 4 and 5 (*/etc/rc4.d/K??nginx* and */etc/rc5.d/K??nginx*) file. This is because *update-rc.d* has been designed in a way to ignore creating or updating any symlinks, if something like */etc/rc?.d/[SK]??name* (where ‘name’ is `nginx`, in our case) already exists. This makes sure that a package update never changes any existing configuration. For details, read the man page of *update-rc.d*.

On RHEL and derivatives:

```
# /sbin/service nginx stop
# /sbin/chkconfig nginx remove
```

Starting the cluster

Now that we are all set to start the cluster, execute the following commands as per your distribution:

```
# invoke-rc.d heartbeat start ## on Debian
or:
# /sbin/service heartbeat start ## on RHEL
```

Our cluster will be up and running in a short time. You can check the availability of the service address and running `nginx` instance. In order to test the cluster failover and fallback, you can start playing with various options like unplugging the network cable from the back of the primary node, switching it off physically, etc.


I will leave testing in your capable hands.

Further tuning

I would strongly recommend that you set up the cluster to use the *ipfail* plug-in to ensure proper failover, if a network

problem occurs. Sometimes, in a large environment where the nodes of the clusters are far apart, the Heartbeat link typically stays alive, and so the cluster thinks that it is functioning in a proper manner. Due to the failure of a switch or router to which the primary node is connected, the resource on the cluster as a whole will not be available to users. In such cases, if you have configured *ipfail* in your cluster, then Heartbeat on all nodes can continuously monitor whether they can reach one of the resources on the network (typically, a switch or router). This resource should not be a member of the cluster. If Heartbeat detects *ping* failures then it can directly go and query if the other node has also detected the failures. If the other node reports that the connectivity is okay, then the services are failed back to that node.

Path to better redundancy

So over the last two articles, we have configured a load-balanced, highly-available cluster of Web servers. There are still many shortcomings in achieving the redundancy for a Web infrastructure. One of these issues is the data centre redundancy. In the next article in this series, we will be looking at setting up a redundant data centre for this environment, as well as at networking set-ups like DNS and firewalls, and the positioning of various components. **END** 

References

- Nginx Home Page: nginx.net
- The HA Linux Project: www.linux-ha.org
- Beowulf Clusters: www.beowulf.org
- Sun Grid Engine: gridengine.sunsource.net
- Wikipedia—Computer Cluster: en.wikipedia.org/wiki/Computer_cluster
- Wikipedia—High-Availability Cluster: en.wikipedia.org/wiki/High-availability_cluster
- Wikipedia—Load Balancing (Computing): [en.wikipedia.org/wiki/Load_balancing_\(computing\)](http://en.wikipedia.org/wiki/Load_balancing_(computing))
- Wikipedia—Beowulf (Computing): [en.wikipedia.org/wiki/Beowulf_\(computing\)](http://en.wikipedia.org/wiki/Beowulf_(computing))
- Wikipedia—Grid Computing: en.wikipedia.org/wiki/Grid_computing
- Wikipedia—Sun Grid Engine: en.wikipedia.org/wiki/Sun_Grid_Engine
- Loadmaster LM-1500: www.kemptechnologies.com/load-balancer-1500.shtml
- Wikipedia - Multicast: en.wikipedia.org/wiki/Multicast
- Wikipedia - Unicast: en.wikipedia.org/wiki/Unicast
- Microsoft KB - Differences between Unicast and Multicast: support.microsoft.com/kb/291786
- Linux Advanced Routing and Traffic Control: lartc.org

By: Ajitabh Pandey

The author has more than 13 years of diversified IT industry experience in training, support and consulting. You can know more about him at ajitabhpandey.info and shoot him an e-mail at [ajitabhpandey \[at\] ajitabhpandey \[dot\] info](mailto:ajitabhpandey[at]ajitabhpandey[dot]info)

LINUX ForYou Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developets or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools





Firewalls, Port Forwarding, NAT, DHCP and TFTP

Last month, we built a server using off-the-shelf hardware. This time, let's set up some essential server services.

Welcome back! For all intents and purposes, we are going to pamper the expensive machine that we built last month. This time, we will look at firewalling, ADSL Router port forwarding, DHCP and TFTP.

Firewalls

Let me break the bad news: There's no firewall software for Linux.

Surprised? Did I just say that all the people who run Linux servers, including the guys at GNU, have no firewall? Of course not. Well, technically, yes, they have no firewall software, but they have a firewall.

Let me explain. All network traffic in

a Linux box is intercepted by the kernel. No direct access is allowed. So, inside the kernel itself, an entire firewall is implemented. We know of this firewall as *iptables*. The *route* command that we executed last month was a part of it.

This *iptables* cannot be called a firewall—it's just a set of rules according to which network traffic is handled. The *route* command that we executed last month added a single default route that all network traffic directed towards the Internet, should take. Rules like this can be used to block viral traffic, or better still, accumulate the viral data in a file for inspection later (that is very complicated, though).

There are several different software

available that act as GUIs for *iptables*, like XFWall. You can use these tools to control your network traffic. XFWall is a very good piece of software, but not entirely documented. Another such software is Firewall Builder, which is also very good (some say the best). However, let's not forget that firewall configuration is unique for every network.

There is another simple, but expensive solution. It constitutes buying an ADSL router (modem), which has in-built firewalling capabilities. My aim was to make the server simple to maintain, and that involves not writing long scripts. If you have a Type-1 ADSL modem (1 port each for power, a phone-line and the network, respectively), then a separate router would be required anyway, because you require port forwarding, which is not present in Type 1 routers. If you go for this, I'd suggest one from Linksys (Cisco)--they're quite easily available.

Network Address Translation

NAT, or Network Address Translation, is a protocol used to forward IP packets from one interface to another. This is a bit different from bridging. Anyway, NAT is essential if we want to browse the Internet from a client. To enable it, execute the following:

```
# echo 1 > /proc/sys/net/ipv4/ip_forward
# iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
# iptables -A FORWARD -i eth0 -o eth1 -m state \
--state RELATED,ESTABLISHED -j ACCEPT
#: iptables -A FORWARD -i eth1 -o eth0 -j ACCEPT
```

In fact, it's better you add these lines to the end of your */etc/rc.local* file to set up NAT every time your server boots.



Note: If the system cannot find the *iptables* command, add an */sbin/* before it. If that too doesn't work, you need to go for an *apt-get install*.

Port forwarding

This is another feature about which I cannot help you much. Port forwarding configuration is different for every modem. Anyway, the instructions here are a bit more general.

Only the ADSL modem can be seen by people on the Internet, not your server. Because of this, all traffic stops at your modem. Port forwarding is enabled to forward all data packets from the modem to the respective port on your computer. If an HTTP request is sent to port 80 of the modem, the request is forwarded to port 80 on your server, on which (hopefully) your HTTP server is running.

You can avoid port forwarding altogether if you do not have an always-on connection. I recommend you go for it, but you don't have to. Make a call to your Internet service provider requesting your modem to be switched back to bridged mode. Use the static IP address you obtained from your ISP here. Now execute *pppoeconf* from your server. Configure the connection. When you are done, run *pon dsl-provider* to bring up the Internet connection. That's it!

DHCP, PXE, DNS and TFTP

Now we begin with the juicy part: setting up the server. Before you start, make sure a client is connected to your server. For simplicity, go for a Linux client.

Issue this command: *sudo apt-get install dnsmasq*. This will install an all-in-one DNS, DHCP and TFTP server on your computer.

Before we start with DHCP, we need to configure the DNS forwarder. The configuration for DNSMASQ is stored in the */etc/dnsmasq.conf* file. We need to edit it. But first, execute these commands:

```
#: mkdir /etc/dnsmasq
#: mv /etc/resolv.conf /etc/dnsmasq/resolv.upstream
#: echo "nameserver 127.0.0.1" > /etc/resolv.conf
```

This moves the previous */etc/resolv.conf* file to */etc/dnsmasq.upstream* and creates a new *resolv.conf* that references the locally running DNSMASQ. Right now, Internet browsing should not be working.

Open the */etc/dnsmasq.conf* file for editing in a text editor and set the following parameters (uncomment them if necessary):

1. First of all, set it to listen on 'loopback' and 'eth1'. Do this by setting two interface lines:

Enterprise Mail Server, Linux SBS Server,
Anti SPAM, Antivirus and HTTP Filtering.

Bandwidth Management, Internet Access Control,
Content Filtering, Web Access Reporting.

Customised Linux Product Development.

TechnoInfotech™

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

- ```
except-interface eth0
except-interface ppp0
```
2. Set DNSMASQ to reference */etc/resolv.upstream* to get the list of upstream name servers:
 

```
resolv-file=/etc/dnsmasq/resolv.upstream
```
  3. Increase the cache size to 1024:
 

```
cache-size=1024
```
  4. Set the domain name to anything you like:
 

```
domain=anything-you-like.local
```
  5. Now the main part: enable DHCP and set the lease time to 'infinite':
 

```
dhcp-range=192.168.1.10,192.168.1.254,255.255.255.0,infinite
```
  6. Set up a Microsoft Windows hack to release DHCP leases on shutdown:
 

```
dhcp-option=vendor:MSFT,2,11
```
  7. Set up the PXE server:
 

```
dhcp-boot=pxelinux.0
enable-tftp
tftp-root=/tftpboot
```
  8. Set some miscellaneous options:
 

```
dhcp-leasefile=/etc/dnsmasq/dnsmasq.leases
dhcp-script=/bin/echo
log-queries
log-dhcp
```
  9. Disable authoritative DHCP: Comment the line, that is put a # sign at the beginning of the line that says "dhcp-authoritative".

Right now, the configuration is a bit broken. Create the directories */tftpboot* and */tftpboot/pxelinux.cfg*.

Download the latest *syslinux tar.bz2* file from [www.kernel.org/pub/linux/utils/boot/syslinux/Testing/](http://www.kernel.org/pub/linux/utils/boot/syslinux/Testing/). Don't worry, 'Testing' releases aren't that broken. Extract it and copy the files *core/pxelinux.0* and *com32/menu/menu.c32* to */tftpboot*.

As an exercise, we are going to download SystemRescueCD and set it up to boot over the network. Download the latest ISO from [nchc.dl.sourceforge.net/sourceforge/systemrescuecd/systemrescuecd-x86-1.1.4.iso](http://nchc.dl.sourceforge.net/sourceforge/systemrescuecd/systemrescuecd-x86-1.1.4.iso), mount it and copy files *rescuecd*, *initram.igz*, *sysrcd.dat*, and *sysrcd.md5* to */tftpboot/sysrcd*.

Copy the following lines to the */tftpboot/pxelinux.cfg/default* file:

```
default menu.c32
prompt 0
```

```
menu title PXE Boot Menu
```

```
label sysrcd
menu label Boot SystemRescueCD
kernel sysrcd/rescuecd
append initrd=sysrcd/initram.igz setkmap=us vga=791 bootftpt=tftp://192.168.1.1/sysrcd/sysrcd.dat
```

Disable PEERDNS if you are using PPPOECONF by commenting (#) the "usepeerdns" line in the */etc/ppp/peers/dsl-provider* file.

Now restart the dnsmasq server:

```
/etc/init.d/dnsmasq restart
```

That should do it. Any machine that accesses PXE when booting will boot into the PXE Bootloader installed on your server. Find out from the motherboard manual of your client how to boot from the network. When you do so, you will be presented with a menu with a single entry.


Press **Enter** now. The kernels, rootfs and the datfile will be downloaded through TFTP and the kernel will be executed. Thus SysRCD will boot up on the client. It will get all its network configuration from DNSMASQ (DHCP).

## What next?

Try browsing the Internet from your newly PXE-booted PC. If all goes well, you should be able to.



**Tips:** You don't need to shut down or restart your server at all from now on. Keep it as it is, and try setting a Guinness Book World Record. On a more serious note, all maintenance, installations and removals will be performed dynamically.

Coming up next, we'll set up two Web server instances on the same server: one to serve your Intranet and one to serve the public through the Internet. Till then, see you! **END** 

**By: Boudhayan Gupta**

The author is a 14-year-old student studying in Class 8. He is a logician (as opposed to a magician), a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++, Python and Assembly (NASM Syntax).

**Got A LUG?**  
Visit [www.openitis.com](http://www.openitis.com)





S.G. Ganesh

# How to Detect Integer Overflow

Integer overflows often result in nasty bugs. In this column, we'll look at some techniques to detect an overflow before it occurs.

Integer overflow happens because computers use fixed width to represent integers. So which are the operations that result in overflow? Bitwise and logical operations cannot overflow, while cast and arithmetic operations can. For example, ++ and += operators can overflow, whereas && or & operators (or even << and >> operators) cannot.

Regarding arithmetic operators, it is obvious that operations like addition, subtraction and multiplication can overflow.

How about operations like (unary) negation, division and mod (remainder)? For unary negation, `-MIN_INT` is equal to `MIN_INT` (and not `MAX_INT`), so it overflows. Following the same logic, division overflows for the expression `(MIN_INT / -1)`. How about a mod operation? It does not overflow. The only possible overflow case `(MIN_INT % -1)` is equal to 0 (verify this yourself—the formula for % operator is "`a % b = a - ((a / b) * b)`").

Let us focus on addition. For the statement "`int k = (i + j);`":

- (1) If *i* and *j* are of different signs, it cannot overflow.
- (2) If *i* and *j* are of same signs (- or +), it can overflow.
- (3) If *i* and *j* are positive integers, then their sign bit is zero. If *k* is negative, it means its sign bit is 1—it indicates the value of (*i* + *j*) is too large to represent in *k*, so it overflows.
- (4) If *i* and *j* are negative integers, then their sign bit is one. If *k* is positive, it means its sign bit is 0—it indicates that the value of (*i* + *j*) is too small to represent in *k*, so it overflows.

To check for overflow, we have to provide checks for conditions (3) and (4). Here is the straightforward conversion of these two statements into code. The function `isSafeToAdd` returns true or false after checking for overflow.

```

/* Is it safe to add i and j without overflow?
Return value 1 indicates there is no overflow;
else it is overflow and not safe to add i and j */
int isSafeToAdd(int i, int j) {
 if((i < 0 && j < 0) && k >= 0 ||
 (i > 0 && j > 0) && k <= 0)
 return 0;
 return 1; // no overflow - safe to add i and j
}

```

Well, this does the work, but is inefficient. Can it be improved? Let us go back and see what *i* + *j* is, when it overflows.

If `((i + j) > INT_MAX)` or if `((i + j) < INT_MIN)`, it overflows. But if we translate this condition directly into code, it will not work:

```

if (((i + j) > INT_MAX) || ((i + j) < INT_MIN))
 return 0; // wrong implementation

```

Why? Because (*i* + *j*) overflows, and when its result is stored, it can never be greater than `INT_MAX` or less than `INT_MIN`. That's precisely the condition (overflow) we want to detect, so it won't work.

How about modifying the checking expression? Instead of `((i + j) > INT_MAX)`, we can check the condition `(i > INT_MAX - j)` by moving *j* to the RHS of the expression. So, the condition in `isSafeToAdd` can be rewritten as:

```

if((i > INT_MAX - j) || (i < INT_MIN - j))
 return 0;

```


That works! But can we simplify it further? From condition (2), we know that for an overflow to occur, the signs of *i* and *j* should be different. If you notice the conditions in (3) and (4), the sign bit of the result (*k*) is different from (*i* and *j*). Does this strike you as the check that the ^ operator can be used? How about this check:

```

int k = (i + j);
if(((i ^ k) & (j ^ k)) < 0)
 return 0;

```

Let us check it. Assume that *i* and *j* are positive values and when it overflows, the result *k* will be negative. Now the condition `(i ^ k)` will be a negative value—the sign bit of *i* is 0 and the sign bit of *k* is 1; so ^ of the sign bit will be 1 and hence the value of the expression `(i ^ k)` is negative. So is the case for `(j ^ k)` and when the & of two values is negative; hence, the condition check with < 0 becomes true when there is overflow. When *i* and *j* are negative and *k* is positive, the condition again is < 0 (following the same logic described above).

So, yes, this also works! Though the *if* condition is not very easy to understand, it is correct and is also an efficient solution! 

## About the author:

**S G Ganesh** is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at [sgganesh@gmail.com](mailto:sgganesh@gmail.com).



# Flirt with Perl



Five rules you should follow.

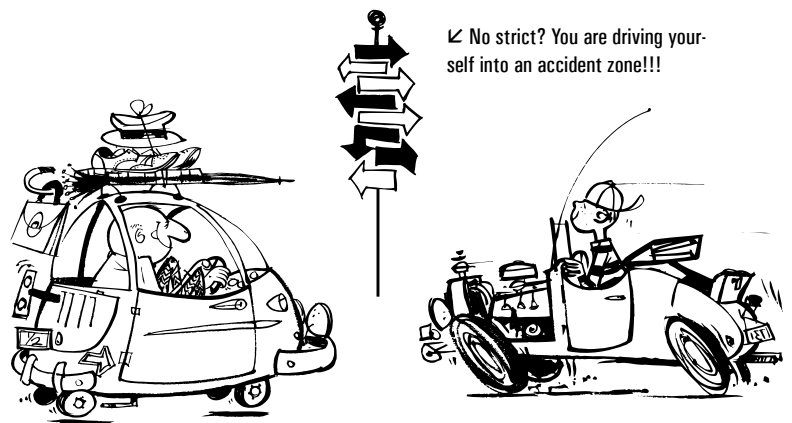
As Damian Conway, a prominent member of the Perl community, once said, “Always code assuming that the guy who’ll end up maintaining your code will be a violent psychopath who knows where you live.” Writing and maintaining code can be one of the most miserable jobs if you don’t follow a certain discipline in your program. At times, the way you implement a solution really matters. Most of us are just bothered about *what* to implement and are completely blind to *how*. This practice can turn a small irritant into a major headache as the volume of code grows. Here are some capsule-sized best practices to avoid such situations.

## ① ‘Strict’ures

Assume a situation in which you are buried in a program of thousands of lines. Somewhere you have misspelled the

name of a variable, say, instead of typing `@friends_list` you punched in `@freinds_list`. You may have no idea that the code is spitting out a bad result because of this silly spelling error. It’s not an easy task to figure out such bugs. Sometimes you really have to struggle for hours to check the databases, file buffers or even the whole logic to fix this naughty bug.

This situation can be easily avoided by using the pragma—*strict*. It forces the programmer to declare all variables as a



package or as lexically scoped before it is used. In fact, every program should start with a *'use strict'* and it is very important that the variables are declared in the smallest possible scope, so as to minimise the 'surprising' outputs.

## ② Modularising

Impatience is an integral characteristic of every programmer. Whenever some task is assigned, we dive into implementing the requirement and start coding pages and pages in a single block/routine. This is a bad habit. It is always worth spending a good amount of time in planning and breaking down your task into smaller atomic routines and giving it a proper name, because *that* is the most crucial part.

↘ Divide and rule



Half your work is done, if this is done perfectly. Writing a single function for eating, exercising and sleeping is also a bad habit. Each piece of code should do one thing and do it well. Furthermore, this practice will make the program more maintainable, and make the debugging and testing much easier. A hash of named arguments are much better than a simple array if you have more than three parameters to be passed to a subroutine. It replaces the need to remember the order of the parameters you pass to the routines.

## ③ Self-documentation

This is nothing but making the code easily readable. It can be achieved by giving the variables, functions, files, etc, a meaningful name. Don't create confusion while naming variables. The name should reflect what it holds.

↙ Clarity makes life easier



This will increase the clarity in your program, which eventually reduces confusion. Since the errors are more obvious, you are less likely to make mistakes. The basic idea of a language, be it spoken or programming, is to communicate. So make your program speak for itself.

## ④ A bit of discipline

Presentation has an enormous impact. Each of the nested blocks should be properly intended. If you don't want to spend time on these cosmetic touches, you can use *Perl::Tidy* which can be downloaded from CPAN. It parses and beautifies your code. This can be even configured according to the style you prefer. Another point on this topic is the naming of your variables and routines. All the names should be largely self-explanatory. Code should read like prose and not like a puzzle.



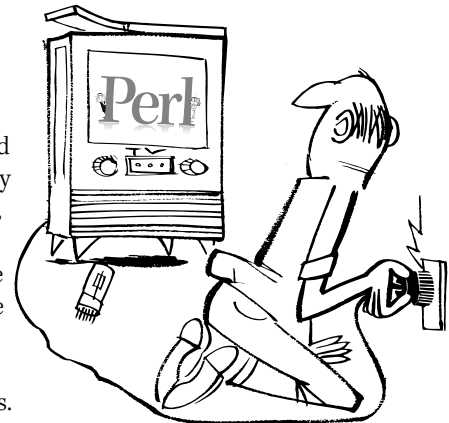
↗ Cosmetic touches are needed

## ⑤ API/module designing

This is all about the designing of the programmer's interface. It requires both experience and creativity. An improper API design will end up in unexpected performance degradation and reduce the maintainability of the system. One good practice of API designing is to write the sample codes that will use the module before writing the module itself. This will help in figuring out how the module should work.

Some of the characteristics of a good API are: easy to use, easy to learn, easy to extend, hard to misuse, easy to read and maintain code that uses it, appropriate to audience and sufficiently powerful to satisfy the requirements.

Starting an API design with a small spec will be ideal, so that it's easy to modify. I strongly recommend reading 'Perl Best Practices' by Damian Conway before a single line of code is written. **END** 



↗ Plug and play

By: Arshad Mohamed

The author is a Web2.0 developer currently working with Genseq, Malaysia as a developer for an online Health Reporting System. He has worked over four years as programmer/developer, and mostly uses Perl. His other expertise includes Ajax, MySQL, Linux, Apache and JavaScript.

# Lynx

## Old, But Still Fresh

A CLI-based browser? Whatever for? Are you still in the early 90s? You may pose all these questions, but the truth is that Lynx, a CLI-based browser, is the favourite of many.

It was once the trusted friend of the visually-impaired—thanks to its text-to-speech friendly interface. But with the advent of better screen readers, Lynx lost some of its regular users (even as some got befuddled!).

You may continue to wonder why people are still using it. Before I elaborate on that, let me show some of its features. In this browser, you can click on a new link by highlighting a chosen link and selecting it. In one of my customised Lynx browsers, I have the freedom to enter the number of the link, as all the links are identified by number.

Support for SSL and many HTML features has been added in the recent versions. Here, the tables in a Web page are linearised. And, just like in Firefox, you have the freedom to view frames as separate pages (in fact, all the frames are identified by names). I am sure that if you haven't been exposed to Lynx, you must be thinking about the non-text content. Lynx can handle these contents by aptly launching an external program for viewing the respective contents—say an image viewer or video player.

Lynx was the brainchild of Lou Montulli, Michael Grobe and Charles Rezac of the University of Kansas (and Thomas Dickey maintains the package now). They brought out the browser way back in 1992. Though it was originally conceived for UNIX and VMS, it is still a popular console-based browser on Linux and is available along with many distros. Figure 1 shows Lynx in Kubuntu.

All the recent versions even run on Windows and Mac OSX (but for Mac there is a 'classical version'

available—MacLynx!). Please refer to the box for the complete list of platforms on which it has been tested.

### Why Lynx?

Now I shall provide you with the reasons why people are still after *this* open-source, text-only Web browser and

#### Lynx has been tested in:

- AIX 3.2.5 (cc w/ curses)
- BeOS 4.5 (GCC w/ ncurses)
- CLIX (cc w/ curses & ncurses)
- DGUX
- Digital Unix 3.2C and 4.0 (GCC & cc w/ curses, ncurses & slang)
- FreeBSD 2.1.5, 3.1 (GCC 2.6.3 w/ curses & ncurses)
- HP-UX (K&R and ANSI cc, GCC w/ curses, ncurses & slang)
- IRIX 5.2 and 6.2 (cc & GCC w/ curses, ncurses & slang)
- Linux 2.0.0 (GCC 2.7.2 w/ curses, ncurses & slang)
- MkLinux 2.1.5 (GCC 2.7.2.1)
- NetBSD
- NEXTSTEP 3.3 (GCC 2.7.2.3 w/ curses)
- OS/2 EMX 0.9c (ncurses)
- SCO OpenServer (cc w/ curses)
- Solaris 2.5, 2.6 & 2.7 (cc & GCC w/ curses, ncurses & slang)
- SunOS 4.1 (cc w/ curses, gcc w/ ncurses & slang)
- OS390 and BS2000

```

File Edit View Terminal Tabs Help
Welcome to Kubuntu 8.04, Hardy Heron! (p1 of 2)

Kubuntu is a free, user-friendly operating system based on the K
Desktop Environment and on the award winning Ubuntu operating system.
With a biannual release cycle and at least 18 months of free security
updates for each release, it is the secure, stable computing
environment you've been waiting for. Its community-driven development
and pervasive availability are grounded in the concept of "humanity to
others."

Kubuntu is not possible without its community, because its projects are
designed, created, and implemented by a diverse community of users.
Individuals and teams provide the code, artwork, documentation, and
technical assistance in addition to promoting Kubuntu to a wider
audience. To read more about how to contribute please view the Helping
Kubuntu page.

The Kubuntu and Ubuntu projects are sponsored by Canonical Ltd, the
leading global provider of free Linux software. Canonical provides full
commercial support for Kubuntu and Ubuntu.

Hit H to open:
Arrow keys: Up and Down to move. Right to follow a link; Left to go back.
Help Options Print G/o Main screen Q/uit /w/search [delete]=history list

```

Figure 1: Accessing [kubuntu.org](http://kubuntu.org) using Lynx from the terminal

Gopher (which is a distributed document search and retrieval network protocol) client. The first reason is that it is quite good when it comes to testing websites. Lynx tries the usability of Web pages in older browsers (see Figure 2).

It is still considered an effective mode to browse the modern Web space. If you have trouble with regard to your low bandwidth or older computer hardware, it's worth giving Lynx a try. The 'speed benefits' associated with Lynx further makes it lucrative to many.

Webmasters and SEOs ought to look at how their websites appear from the eyes of a spider. Many of them use Lynx for that! In the Web sphere, there are even many 'Lynx viewers' that allow them to have a glance at the Web pages, using emulators (rather than the original Lynx). Figure 3 shows one such 'view'. This further facilitates the Web master to figure out (in a critical sense) whether his site is accessible to the visually impaired.

Using Lynx, you can verify whether a website is crawled correctly by a search engine. Web pages written with Lynx in mind often receive good page ranks, as robots, index or abstract generation tools can leniently imbibe and extract data from the page.

Further, you can easily modify it (as it is in ISO C) to suit your needs.

## So what's the downside?

Of course, there are a few negatives associated with the current form of Lynx. The most prominent is the bad HTML handling when it comes to forms. Lynx's comment handling is also somewhat poor, especially when it meddles with a slightly incorrect HTML format. Even the cookie implementation doesn't seem to be perfect.

Redirecting POST content is another facet that needs improvisation. Lynx may ask you:

"wWw: Redirection for POST content. Proceed (y/n)? "

If your response is 'No', the server may get perplexed and you may get an error.

But Lynx is still the favourite browser of many.

```

File Edit View Terminal Tabs Help
The Analyst Magazine (p1 of 3)

[Top-Logo.gif]

HOME CURRENT ISSUE NEWS ARCHIVES EDITORIAL
BOARD ABOUT CONTACT

Rousing from the perceptions of GNU and Free Software Foundation, The
Analyst magazine was launched so as to build an environment where one
could find quality analysis reports of current events and other
educational resources for free. We believe that, just like softwares,
all educational resources should obtainable without any constraints.
This magazine is expressly intended for the benefit of people who wish
to have a glance at the world affairs from an unbiased perspective. The
Editorial Board has taken utmost exertion to guarantee the authenticity
and quality of the articles. Unlike projects like Wikipedia, The
Analyst Magazine writers in order to uphold the quality of the
articles. We also tag on to the copyleft policy. Our sources of revenue
and other policy matters are well elucidated in the 'Editorial Board'
Page. The whole revenue from the site is used for maintaining the site
and for paying for writers.

-- press space for next page --
Arrow keys: Up and Down to move. Right to follow a link; Left to go back.
Help Options Print G/o Main screen Q/uit /w/search [delete]=history list

```

Figure 2: Home page of The Analyst magazine in Lynx

```

[Top-Logo.gif]

HOME [1]CURRENT ISSUE [2]NEWS [3]ARCHIVES
[4]EDITORIAL BOARD [5]ABOUT [6]CONTACT

Rousing from the perceptions of GNU and Free Software Foundation, The
Analyst magazine was launched so as to build an environment where one
could find quality analysis reports of current events and other
educational resources for free. We believe that, just like softwares,
all educational resources should obtainable without any constraints.
This magazine is expressly intended for the benefit of people who wish
to have a glance at the world affairs from an unbiased perspective.
The Editorial Board has taken utmost exertion to guarantee the
authenticity and quality of the articles. Unlike projects like
Wikipedia, The Analyst Magazine writers in order to uphold the
quality of the articles. We also tag on to the copyleft policy. Our
sources of revenue and other policy matters are well elucidated in the
'Editorial Board' Page. The whole revenue from the site is used for
maintaining the site and for paying for writers.

[7]Dr. Stallman's comments on The Analyst Magazine

[8]READ RELEASE UPDATES - AND KEEP UPDATED

[9]SPECIAL REPORTS
[10]RESOURCES
[11]VIDEOS
[12]PARTNERS
[13]FREELANCERS
[14]COMPETITIONS
[15]ADVERTISE
[16]LINKS
[17]SEARCH
[18]FORUMS Make payments with PayPal - it's fast, free and secure!

st Popular Articles - All time

```

Figure 3: Home page of The Analyst magazine in a Lynx viewer

You can easily find many volunteers who can help you while meddling with its code. Give it a try, and you can experience the advantages of a CLI-based browser!



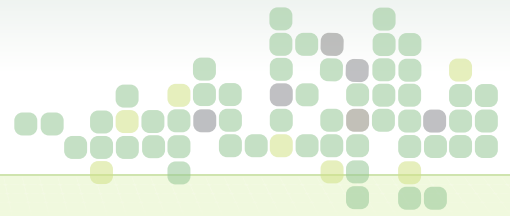
### Resources

- Lynx home page: [lynx.isc.org](http://lynx.isc.org)
- Lynx viewer: [www.yellowpipe.com/yis/tools/lynx/lynx\\_viewer.php](http://www.yellowpipe.com/yis/tools/lynx/lynx_viewer.php)
- Extremely Lynx: [linux4u.jinr.ru/usoft/WWW/www\\_crl.com/subir/lynx.html](http://linux4u.jinr.ru/usoft/WWW/www_crl.com/subir/lynx.html)
- See the web like a crawler: [seebot.org](http://seebot.org)
- An Early History of Lynx: [people.cc.ku.edu/~grobe/early-lynx.html](http://people.cc.ku.edu/~grobe/early-lynx.html)

### By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at [www.aasisvinayak.com](http://www.aasisvinayak.com)





# Industry News

## Brazil witnesses world's largest desktop Linux deployment

Userful and ThinNetworks have been selected to supply 3,56,800 virtualised desktops to schools in all of Brazil's 5,560 municipalities. According to the duo, it is a historical achievement of being the world's largest ever virtual desktop deployment; the world's largest ever desktop Linux deployment; and a new record for low-cost PCs -- with the PCs sharing hardware and the software, hence costing less than \$50 per seat.

"Userful is very happy to have been selected to participate in this historic opportunity to help millions of children get the computer education they need in a sustainable way," said Tim Griffin, president, Userful.

In combination with hardware from ThinNetworks, Userful Multiplier is also the lowest-cost desktop virtualisation solution in the market. Userful offers the features of a full PC, including high performance video for less than \$50 per additional seat in large deployments (not including monitors and keyboards), and uses standard PC hardware including additional low-cost video cards and USB/2-way-audio hubs from ThinNetworks.

"This deployment alone saves more than 170,000 tons of CO2 emissions annually, the same as taking 28,000 cars off the road, or planting 41,000 acres of trees," said Sean Rousseau, marketing manager, Userful. "Turning 1 computer into 10 reduces computer hardware waste by up to 80 per cent, further decreasing its environmental footprint."

The project comprises three phases, the first of which consist of 18,750 workstations in rural schools that have already been installed and are in operation.



## Wind River becomes LiMo's systems integrator

LiMo Foundation has selected Wind River, a device software optimisation (DSO) company, as the systems integrator to deliver the common infrastructure, tools, testing and integration services for the LiMo platform. A board member of LiMo Foundation, Wind River will now provide technology and services to combine, harden and validate code contributions of LiMo members through a Common Working Environment (CWE).

"As the mobile industry now breaks out of its traditional, controlled development environments and embraces collaborative approaches that unlock innovation, operators and device manufacturers are turning to vendors they can trust to guide them through the evolving mobile software landscape," said Morgan Gillis, executive director, LiMo Foundation. "Through its selection of Wind River as a company with a legacy of world-class embedded software, LiMo expects to be able to accelerate its goal of reducing Linux fragmentation while intensifying platform development through the rapid adoption of member contributions."

In addition to the integration of member code, Wind River expects to enhance the efficiency of the contribution process for members through collaboration tools and enhanced processes. This will become increasingly important as more strategic contributions require integration into LiMo reference platforms.

## Sun launches 'Open Innovation Portal'

Reaffirming its commitment to open source as a catalyst to encourage innovation in India, Sun Microsystems has launched the Open Innovation Portal at the Centre for Excellence in e-governance, Department of Management Studies, IIT Delhi with JNU and Knowledge Commons as collaborators.

The portal [[www.innovationcommons.org](http://www.innovationcommons.org)] was launched by Joe Hartley, Sun's VP for Global Education, Government and Healthcare, and Prof S. S. Yadav, head of the Department of Management Studies, IIT Delhi. The objective behind the launch is to foster the development of participative innovation in society, and to help transition the economy into an innovation economy. The portal will allow innovators from the science and student community and from all walks of life to publish their innovations with no IPR encumbrances, for the whole world to benefit. It will also help others to add to any open innovation that has been published online, thus enabling the process of participative innovation.

Speaking at the inauguration, Hartley said, "The Innovation Portal therefore allows for collaboration between members of a society

to foster innovation. This is in line with Sun's vision of the Participation Age in which, the company believes, the world has



entered a new era where dramatically lowered barriers to entry, plummeting device prices, and near-universal connectivity are driving a new round of network participation."

## LG, Intel collaborate on future MIDs

LG Electronics and Intel have announced a collaboration for mobile Internet devices (MIDs). This is based on Intel's MID hardware platform, codenamed Moorestown, and Linux-based Moblin v2.0 software platform. The LG device is expected to be one of the first Moorestown designs to enter the market.

LG and Intel's common goal is to unleash rich Internet experiences across a range of mobile devices while delivering the functionality of today's high-end smart phones. The collaboration on the new design extends the close working relationship the two companies have enjoyed across their respective mobile product lines, which now spans the notebook, netbook and MID categories.

"The MID segment will drive growth at LG Electronics. We chose Intel's next-generation Moorestown platform and Moblin-based OS to pursue this segment because of the high performance and Internet compatibility this brings to our service provider customers," said Jung Jun Lee, executive vice president of LG Electronics and head of its Mobile Communications Business Division. "The collaboration with Intel on the MID platform has been valuable and further extends our longstanding relationship. Our efforts are well on track and we look forward to bringing the MID to market."

## Red Hat, Microsoft in virtualisation inter-op pact

Red Hat and Microsoft customers will now be able to run Microsoft Windows Server and Red Hat Enterprise Linux virtual servers on either host environment, with configurations that will be tested and supported by both virtualisation and operating system companies. Red Hat, in response to strong customer demand, has signed reciprocal agreements with Microsoft to enable increased interoperability for the companies' virtualisation platforms.

Each company will join the other's virtualisation validation/certification programme and will provide coordinated technical support for their mutual server virtualisation customers. The reciprocal validations will allow customers to deploy heterogeneous, virtualised Red Hat and Microsoft solutions with confidence, said sources from the newly signed partnership. "The world of IT today is a mixture of virtualised and non-virtualised environments. Red Hat is looking to help our customers extend more rapidly into virtualised environments, including mixed Red Hat Enterprise Linux and Windows Server environments," said Mike Evans, vice president, corporate development, Red Hat.

The key components of the reciprocal agreements are: Red Hat will validate Windows Server guests to be supported on Red Hat Enterprise virtualisation technologies; Microsoft will validate Red Hat Enterprise Linux server guests to be supported on Windows Server Hyper-V and Microsoft Hyper-V Server; and once each company completes testing, customers with valid support agreements will receive coordinated technical support to run the Windows Server OS virtualised on Red Hat Enterprise Linux, and to run Red Hat Enterprise Linux virtualised on Windows Server Hyper-V and Microsoft Hyper-V Server. The agreements establish coordinated technical support for Microsoft and Red Hat's mutual customers using server virtualisation, and the activities included in these agreements do not require the sharing of IP. Therefore, the agreements do not include any patent or open source licensing rights, and additionally contain no financial clauses, other than industry-standard certification/validation testing fees.

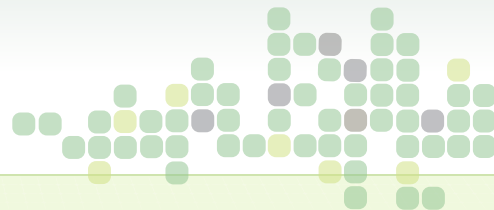


## First open source standard for storage encryption solutions

Sun Microsystems has released the world's first generic communication protocol between a Key Manager and an encrypting device into the open source community. According to Sun, this latest effort in open storage gives customers greater choice, value and flexibility through the resources in open source communities, like the growing storage community within OpenSolaris. The release enables partners to adopt the protocol to securely handle encryption keys without additional licensing. The protocol is implemented as a complete toolkit and is downloadable from the OpenSolaris website [opensolaris.org/os/project/kmsagenttoolkit](http://opensolaris.org/os/project/kmsagenttoolkit).

"Open Storage solutions allow customers to break free from the chains of proprietary hardware and software, and this new protocol extends this lifeline into the expensive and highly fragmented encryption market," said Jason Schaffer, senior director, storage product management, Sun Microsystems. "Open source equals customer value for encryption solutions, and Sun now offers the only solution on the market that works across multiple vendors and suppliers."





## **GCC goes GPLv3**

The Free Software Foundation (FSF), together with the GCC Steering Committee and the Software Freedom Law Centre, announced the release of a new GCC Runtime Library Exception. This licence exception will allow the entire GCC code base to be upgraded to GPLv3, and enable the development of a plug-in framework for GCC. "GCC includes runtime libraries that are automatically built into all the object code that GCC creates," explained Brett Smith, licence compliance engineer at the FSF. "Because we decided a long time ago to allow developers to compile proprietary software with GCC, these libraries have always had licence exceptions. This way, programs that are merely compiled with GCC don't have to be released under the GPL."

The text of the exception is available at [www.fsf.org/licenses/licenses/gcc-exception.html](http://www.fsf.org/licenses/licenses/gcc-exception.html). The FSF has also published a rationale document and FAQ at [www.fsf.org/licenses/licenses/gcc-exception-faq.html](http://www.fsf.org/licenses/licenses/gcc-exception-faq.html) to help users understand the exception better.

## **Azingo to create Vodafone's Linux apps**

Vodafone has selected the open mobile OS company, Azingo, as a partner to develop applications for phones based on the LiMo platform. Both Azingo and Vodafone are core members of the LiMo (Linux Mobile) Foundation, which is an alliance started by Motorola, NEC, NTT DoCoMo, Panasonic Mobile Communications, Samsung Electronics and Vodafone in January 2007. Since then, a number of new members (like the Mozilla Foundation) have joined the alliance.

"We are excited to partner with Azingo to develop cutting-edge applications for our mobile phones based on the LiMo platform," said Guido Arnone, director, terminals technology, Vodafone. "We're looking forward to working with Azingo's agile development teams to develop and deliver innovative communications solutions for our customers."



## **HP bundles Blade PC and Citrix XenDesktop**

HP and Citrix have announced a simplified solution that integrates affordable, high-performance HP Blade PCs and Citrix XenDesktop 3 to help businesses reduce costs and enjoy better manageability, scalability and security than with traditional PCs. "By combining HP blade PCs with XenDesktop, we've created a simple, low-cost virtualisation solution that helps companies efficiently manage and scale their computing environments while delivering the high-performance experience that knowledge and power users demand," said Roberto Moctezuma, vice president and general manager, Desktop Solutions Organization, HP. "Particularly in this challenging economic environment, we see client virtualisation as a cost-efficient alternative for companies needing to economically update and better manage their personal computing infrastructures."

At the core of the solution are the new HP BladeSystem bc2800 Blade PC and HP BladeSystem bc2200 Blade PC, which offer advanced infrastructure control and scalability. These products are combined with Citrix XenDesktop 3 to provide a high-definition user experience and centralised desktop management. The combined offering leverages the power of both the data centre and endpoint devices to significantly reduce desktop TCO.

## **Bank of New Zealand deploys Red Hat on mainframes**

Red Hat has announced that the Bank of New Zealand, a subsidiary of the National Australia Bank Group, has deployed Red Hat Enterprise Linux (RHEL) 5 on IBM System z mainframes to solve environment, space and cost issues related to its data centres. With Red Hat and IBM solutions, Bank of New Zealand has significantly reduced its hardware footprint, power consumption, heat and carbon emissions and costs, including an expected 20 per cent cost reduction over the life of the platform.

"Bank of New Zealand had defined two important goals for the future, both of which relied heavily on IT. The first was for the organisation to become carbon neutral by 2010 and the second was to explore open source opportunities through the adoption of Linux," said Lyle Johnston, infrastructure architect at BNZ. "We also faced the challenge of creating a disaster-recovery solution for our data centres in Auckland, New Zealand and East Melbourne, Australia."

In mid-2007, BNZ began overhauling its mission-critical front-end IT environment, including its Internet banking and bank teller functions, and its middleware layer providing connectivity through to its core back-end data. It migrated its systems to RHEL 5 running under z/VM on the mainframe. Today, BNZ utilises both IBM System z10 and z9 systems, exclusively running RHEL 5, to power the bank's customer-facing banking systems, including Internet banking and teller platforms. "We have also managed to substantially reduce our front-end power consumption by nearly 40 per cent, which means we are well and truly on our way to becoming carbon-neutral by our target year of 2010," said Johnston.





# We have crossed **2.5 million** hits per month !



Join These Growing Numbers And Benefit...

## electronicsforu.com

A comprehensive portal on electronics

### Popular sections



#### Circuit Lab

Hundreds of circuit diagrams of different projects, for **FREE**



#### Top 100

Most popular technology websites, measured by their hits



#### Events

A comprehensive list of technical events, in India and abroad

Looking for something related to electronics? Try electronicsforu.com. This website has circuit diagrams, technical articles, bulletin board, dedicated index of electronics websites, list of popular websites ranked by hits...you name it, it's got it.

Check it out today!



EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020  
Phone: (011) 26810601-03 Fax: (011) 26817583/26812312 Email: info@efyindia.com





Sandya Mannarswamy

Welcome to another installment of CodeSport. In this month's column, we'll explore the best lower bounds of algorithms to determine whether a given graph is connected or not. We will then discuss the problem of finding the minimum element in a circular sorted linked list, given an arbitrary pointer into the list.

Thanks to all the readers who commented on the problems we discussed in last month's column. Last month's takeaway question was to consider the well-known problem of deciding whether a given graph with  $N$  nodes is connected, and determine its best lower bound. The only question the algorithm can ask the adversary is of the form, "Does an edge exist between Vertex  $u$  and Vertex  $v$ ?" The readers were asked to come up with the best lower bound they could establish for this algorithm, using an adversary argument.

Since none of the solutions I received for this problem were completely correct, I am going to keep this problem open to readers this month also. However, in order to help those who were pretty close to getting the correct answer, I will give you some clues. Before trying to determine the lower bound for the 'graph connectedness' problem, let us first try to solve the problem using an algorithm known to us. Many of the questions on graphs, such as graph connectedness, or the presence of a cycle or path between 2 vertices, can be solved by using a variant of the 'depth first search traversal', which we have discussed in an earlier column.

## How can you use depth first search to determine graph connectedness?

As we know, the depth first traversal is a graph traversal wherein we start visiting the children at each node, following any children that may exist at the current node, till we hit a leaf node. We then traverse back up one level, and restart the visit from there. This is unlike a breadth first search where nodes at each level are visited fully before any nodes at the next level are visited. For each node in the graph, perform a DFS from that node which is marked as 'root' for the current traversal iteration. If all other nodes of the graph cannot be visited during the depth first search

from the current root, then the graph is not connected. What is the complexity of such an algorithm? We know that DFS has a complexity of  $O(V+E)$ , where  $V$  is the number of vertices of the graph and  $E$  is the number of edges of the graph. Since  $E$  can be of the order of  $O(V^2)$ , the complexity is  $O(V^2)$ .

## Lower bound for graph connectedness using adversary argument

Coming back to the question of the best lower bound for graph connectedness using an adversary argument, remember that if we do not probe all the possible edges of the graph using the question, "Is Vertex  $v$  adjacent to Vertex  $u$ ?" it is possible for the adversary to answer all our questions correctly and still leave Vertex  $u$  totally disconnected from Vertex  $v$ , resulting in the graph being disconnected. Therefore, any algorithm that gives the correct answer, must probe edges between each pair of vertices in the graph. Given a graph with  $N$  vertices, what is the maximum number possible for the edges in an undirected graph? The complete graph on ' $N$ ' vertices has the maximum number of edges and it is equal to  $NC_2$ , since an edge connects 2 vertices, and the number of ways of choosing any two vertices out of  $N$  vertices is  $(NC_2)$ . Note that  $NC_2$  stands for choosing a combination of 2 items out of  $N$  items and is given by the formula  $(N*(N-1))/2$ . Hence, any algorithm to determine graph connectedness must examine all  $(NC_2)$  edges. Thus the lower bound of any graph connectedness algorithm cannot be less than  $(NC_2)$ . With this clue, the readers should be able to solve the question of the best lower bound for graph connectedness using the adversary argument.

## The question this month

In this month's column, we will revisit number searching. You are given a circular list of ' $n$ ' numbers and the numbers on the list are strictly increasing. Since it is a circular list, the end of the

list wraps over to the beginning of the list. You are given an arbitrary pointer to an element in the list. You need to find the minimum element in the list. You can make the simplifying assumption that all the elements are distinct.

## Finding the minimum in a circular sorted linked list

The simplest approach is to use linear search. Start traversing the list from the pointer you have been given. If you reach a node whose value is smaller than the previous node value, then you have reached the minimum. Given below is the pseudo-code for this solution:

```
struct node* find_min(struct node* p)
{
 struct node* curr = p;
 struct node* prev = NULL;

 do
 {
 If (prev != NULL)
 {
 prev_val = prev->value;
 curr_val = curr->value;
 if (curr_val < prev_val) //we have found the minimum
 return curr;
 }
 prev = curr ;
 curr = curr->next;
 } while (curr->next != p);

 return p;
}
```

## Complexity of the simple solution

What is the time complexity of the simple solution? Since we are doing a linear search on the linked list, it is  $O(n)$ —where  $n$  is the number of elements on the linked list. We can speed up the solution somewhat by making the pointer move forward by  $2/4/8$ , and if we find that the value of the current node is smaller than the value of the previous node we looked at, we know that the minimum lies between these two and we need to linearly search this range. However, this does not bring down the overall complexity which remains at  $O(n)$ . I leave it to the reader to write the pseudo code for the solution that incorporates this pointer-jumping enhancement.

## How can we further improve the solution?

As you may have noticed by now, we are told that the list of numbers we have is strictly increasing until we reach the end of the list, at which point it wraps over, and we again have a strictly increasing list of numbers from that point. This circular list can be viewed as two list segments, each containing the sorted list of numbers—the first segment starting from the arbitrary pointer we have been given to the end of the linked list, and the second segment starting from

the true beginning of the list to the element just before the arbitrary pointer we have been given. Let us refer to these segments as SEG1 and SEG2. This helps us mull over these segments easily. If SEG2 has zero length, what does it mean?


It means that the arbitrary pointer we have been given is in fact the true beginning of the sorted list and, hence, the minimum is pointed to the pointer 'p' we have been given. If SEG2 has non-zero length, we need to determine the start of SEG2, since that would point to the true beginning of the list and hence will contain the minimum. I leave it to the reader to consider the case when SEG1 has zero length.

Can we reduce the time complexity by employing some form of binary search on the circular linked list? As the reader can see, this is not possible, since binary search requires random access, and given that we have a circular linked list, we cannot use binary search to speed this up. This is a typical example of a case, where due to the limitation of the underlying data structure (the circular linked list, in this case), we cannot employ an algorithm to speed up the solution. Hence, linear search remains the best solution for this case.

Now, if we consider a variant of the above problem wherein we lift the restriction that the list is in fact represented by a circular singly linked list, and allow the list to be represented by a circular doubly linked list, can we do any better? It is easy to show that since a linked list does not support random access, using a doubly linked list does not bring down the asymptotic complexity from  $O(n)$ . Now let us assume that we relax the constraints even further and allow an array to represent the circular list, can we use binary search to improve the  $O(n)$  solution we have achieved? I leave this question to the reader to answer.

## This month's takeaway problem

This month's takeaway problem comes from graph theory. Given a directed graph, a Strongly Connected Component (SCC) is a set of nodes such that every pair of nodes in the SCC is mutually reachable from another. In simple terms, SCC corresponds to reducible loops in the graph. A directed graph can contain multiple strongly connected components. One of the interesting applications of SCCs is to represent the loops in the software programs we write. For example, a *for* loop you write in your 'C' code is internally represented as a SCC when the compiler analyses the function for optimisation opportunities. Can you come up with an algorithm to find all the strongly connected components of the graph?

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me. Feel free to send me your solutions and feedback to sandyasm\_AT\_yahoo\_DOT\_com. Till next month, happy programming! 

### About the author:

**Sandya Mannarswamy.** The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



# A Peek Into the WWW, Courtesy MozillaCamp



Delhi's first unconference on Mozilla technologies was a grand event with about a 100 campers who came together to share some Mozilla love on February 10. It was an event that attracted technologists and students, with Mozilla's Seth Bindernagel and Arun Ranganathan around to discuss the future of the Web.

rganised by the Mozilla Community in Delhi, the event was the result of the efforts put in by its two unorganisers—Mohak Prince, Mozilla campus ambassador for Maharaaj Agrasen Institute of Technology and Kinshuk Sunil, community manager for OSSCube. The event was wholly sponsored by OSSCube and was supported by Routeguru, Innobuzz, Pringoo, Pictualise, BlogAdda, Indyarocks, and *LINUX For You*.

Despite February 10<sup>th</sup> being a Tuesday, a lot of participants, from professionals to students, joined the

unconference. There were a few initial worries with Seth and Arun not being around pre-lunch. Mohak and Kinshuk undertook some open house sessions discussing Twitter, Google Chrome, and the browser wars, among other things. Gaurav Paliwal, a student of IP University, led a technical session on SLIM Server. Manu Goel, a UI designer at Sapient delivered a talk on 'Design Trends of Web 2.0', which led to a healthy discussion on the user experience and user interface.

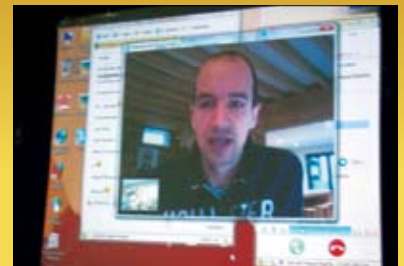
Post lunch, Arun and Seth joined the participants.



Seth Bindernagel shares information on Mozilla's Localisation activities



Arun Ranganathan shares his insight on Web technologies




Pascal Finette talking about 'Concept Series'

After a brief introduction by both of them, Arun put through a Skype call to Pascal Finette of Mozilla Labs, who took the opportunity to walk the participants through the what and how of the Concept Series [[labs.mozilla.com/projects/concept-series](http://labs.mozilla.com/projects/concept-series)]. At the end of the session, the house was opened to some very interesting questions by the participants.

Next, Seth introduced Mozilla to the participants and took them through the efforts Mozilla has been making on the localisation front. He also introduced the audience to Silme [[wiki.braniecki.net/Silme](http://wiki.braniecki.net/Silme)] and walked them through its many uses.

This was followed by a presentation on the evolution of Firefox. This lively session was a contribution by Pictualise. Arun then took control of the house for his talk on the Open Web. He also showcased a number of experimental features coming up in Firefox 3.1 and HTML 5, especially the 'video' tag and Canvas. His talk was filled with a lot of interesting trivia and tidbits on the evolution of the Web, starting from the days of the browser wars.

By 5.30 in the evening, the event wound up, after which many participants thronged to Arun and Seth to discuss and share a number of issues. Subsequently, Arun, Seth, Mohak and I walked to the India Habitat Centre for an interview, which is included in the videos from the event. Have a look at it AT [vimeo.com/album/65868](http://vimeo.com/album/65868), if you don't want to miss anything from the event.

We are grateful to our sponsors and supporters for the help extended to us and to all the participants who helped by contributing to the event. Special mention should be made of the 'Twitterers' who 'live tweeted' the whole event so that those who could not be physically present could still follow everything that happened. To check out on those tweets look for the [#mozcampdel](#) at [search.twitter.com](http://search.twitter.com). **END** 

#### By Kinshuk Sunil

The author is the community manager at OSSCube, and helped organise the MozillaCamp Delhi. He is also an active volunteer for other unconferences in the FOSSverse and an active member of the OSScamp community. You can catch him at [www.kinshuksunil.com](http://www.kinshuksunil.com)



The event registration desk kept busy as attendees poured in

## OSScamp goes to Uttarakhand

OSScamp Panthnagar was organised at College of Technology's (COT) Department of Computer Engineering, on January 31, 2009. Around 250 campers registered at the camp, including 50 enthusiasts from outside Panthnagar. The unconference began with Kinshuk Sunil briefing the audience on open source concepts and software licensing, followed by Sapient's Manu Goel delivering a talk on JavaScript, DOM and AJAX. Next up was SUN's Ajay Ahuja who shared some insights on OpenSolaris and its unique features like, ZFS and DTrace. He also introduced the participants to new technologies like OpenSPARC and SUN's contribution to Linux.

Toshendra Sharma (a final year student, IT, COT, Panthnagar) explored prospects in 3D animation in Open Source with Blender. Another student, Kanika Singhal (third year, computer engineering, COT) explained beginner level PHP coding and PHP's advantages over JavaScript, ASP, etc. Finally, Manu Goel entertained participants' request to talk on CSS and gave live demonstrations that marked the end of the day-one of the event.

Day-two started with an informal introduction of the experts, followed by a talk on networking. TCP/IP, tunnelling, proxy, IP address and the loopholes in network security were discussed thoroughly by Rony Felix of OSSCube.

The second half witnessed talks from COT students. Saurabh Saxena (a final year student) spoke on the subject of advanced PHP. Kartik Asooja (third year), Saurabh Shekar Verma (third year) and Abhinav Pundeer (second year) explored the prospects of OpenGL. An open discussion on Javascript and AJAX by Manu Goel, marked the end of OSScamp Panthnagar Chapter-I. Time to open Chapter-II?

**By: Vidushi Rastogi**, computer engineering student at COT, Panthnagar, and Priyanka Jain of OSSCube.

## Upcoming OSScamps

### Chennai Camp

**Theme:** Everything Open Source

**Audience:** Professionals, Developers, Students, Technology Enthusiasts

**Date:** March 13-14, 2009 with OSI Tech Days

**Venue:** Chennai Trade Center, Chennai

**Details:** [chennai.osscomp.in](http://chennai.osscomp.in)

### Delhi Camp

**Theme:** Everything Open Source

**Audience:** Professionals, Developers, Students, Technology Enthusiasts

**Date:** March 28-29, 2009

**Venue:** Indian Institute of Technology, Delhi

**Details:** [delhi.osscomp.in](http://delhi.osscomp.in)





# Tips & Tricks



## How to split the files

The following is an example of how to use the *split* command on a 600MB *image.iso* file:

```
split -b 200m image.iso
```

It will generate three files, namely xaa, xab and xac, of 200MB each. Afterwards you can use the *cat* command to combine the three to get back the original file, as follows:

```
cat xaa* > new-image.iso
```

—Remin, [remin13@gmail.com](mailto:remin13@gmail.com)



## Interface devices fail to start up at boot time

Open the network card configuration script (it could be */etc/sysconfig/network/ifcfg-ethX* or */etc/sysconfig/network-scripts/ifcfg-ethX*, depending on the distro) in a text editor and add the following line:

```
STARTMODE=auto
```

To manually start it up:

Use */sbin/ifup* and */sbin/ifdown* to start and stop the network interface.

—Dheep Surendran, [bobby\\_dheep@yahoo.com](mailto:bobby_dheep@yahoo.com)



## Lost Bash history

If you have a terminal open in which you're executing certain commands, then open another one and use that for a while. You'll notice this new terminal doesn't remember any of the commands typed in the first one.

In addition, closing the first terminal, and then the

second will overwrite any of the commands typed in the first terminal. This happens because Bash history is only saved when you close the terminal, not after each command. To fix this, add the following lines to your *~/.bashrc* file:

```
shopt -s histappend
```

```
PROMPT_COMMAND='history -a'
```

This will make Bash append an entry in its history after the execution of every command.

—Govindarajalu, [govind.rajalu@gmail.com](mailto:govind.rajalu@gmail.com)



## Prevent users from changing their passwords

Usually */usr/bin/passwd* has the following SUID permission:

```
-r-s--x--x 1 root root 19348 Sep 7 2004 /usr/bin/passwd
```

The numerical value of the file permission translates to 4411.

When a SUID file is executed, the process that runs it is granted access to system resources based on the user who owns the file and not the user who created the process.

So, we need to remove the SUID for that command, so that normal users are denied the privileges of updating the file:

```
chmod u-s /usr/bin/passwd
```

...or:

```
chmod 511 /usr/bin/passwd
```

—Govindarajalu, [govind.rajalu@gmail.com](mailto:govind.rajalu@gmail.com)



## Change the message on your login page

If you need to change the message of your virtual console login screen, edit the `/etc/issue` file and write the message that you want to appear at your login screen.

—Bharat Kumar, [kmr.bharat@gmail.com](mailto:kmr.bharat@gmail.com)



## Find a word

To search for a particular stating word in a file, you can use the `find` command in the following way:

```
find / -type f -exec grep -H 'Suyash' {} \;
```

This command will search for the word `Suyash` in the entire file system. If you only want to search in a particular file/folder, you need to specify the path as the first argument. Note that I have provided `"/` because I wanted to search the entire file system.

—Suyash Jain, [sjain1980@gmail.com](mailto:sjain1980@gmail.com)



## Extract the contents of an RPM

Sometimes we are required to extract the files inside an RPM file instead of installing the RPM. A good example is when we take binaries from one distribution and to use on another distribution, where RPM is not the default package manager. The `rpm2cpio` command comes in handy under these circumstances. For example:

```
$ rpm2cpio coreutils-6.9-2.fc7.i386.rpm | cpio -idv
```

```
./bin/basename
```

```
./bin/cat
```

```
./bin/chgrp
```

```
./bin/chmod
```

```
[...]
```

This command can be used for source RPMs also.

—Yogindar Das Y, [yogindar@yahoo.com](mailto:yogindar@yahoo.com)



## Checking memory and I/O

The `vmstat` utility provides interesting information about processes, memory, I/O and

CPU activity. When you run this utility without any arguments, the output looks similar to the following:

```
procs memory swap io system cpu
rbw swpd free buff cache si so bi bo in cs us sy id
000 8 8412 45956 52820 0 0 0 0 104 11 66 0 33
```

Here, the 'procs' fields show the number of processes:

- Waiting for run time (r)
- Blocked (b)
- Swapped out (w)

The 'memory' fields show the KBs of:

- Swap memory
- Free memory
- Buffered memory
- Cached memory

The 'swap' fields show the KBps of memory:

- Swapped in from disk (si)
- Swapped out to disk (so)

The 'io' fields show the number of blocks per second:

- Sent to block devices (bi)
- Received from block devices (bo)

The 'system' field shows the number of:

- Interrupts per second (in)
- Context switches per second (cs)

The 'cpu' field shows the percentage of total CPU time as:

- User time (us)
- System time (sy)
- Idle (id) time

If you want `vmstat` to update information automatically, you can run it as `vmstat nsec`, where `nsec` is the number of seconds you want it to wait before another update.

—Ashish Kumar, [Suyash Jain](mailto:Suyash Jain)



### Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at [www.linuxforu.com](http://www.linuxforu.com). The sender of each published tip will get an LFY T-shirt.

# A Voyage to the Kernel



Part 10

## Day 9 | Segment 2.3

### Algorithms in cryptography

The study of secret communication systems has lured people of all ages. And the old methods of encrypting messages are quite popular even in literature. But our interest is centred around two aspects-cryptography and cryptanalysis. Cryptography, in plain words, is concerned with the design of secret communications systems, while the latter studies the ways to compromise secret communications systems!

We all know that when a bank upgrades its systems to incorporate IT, it has to make sure that the methods of electronic funds transfer are just as secure as funds transferred by an armoured vehicle.

You might have seen the arithmetic and string-processing algorithms that people employ in this realm, which are what beginners are expected to study.

Cryptanalysis, for sure, can place an incredible strain on the available computational resources. That is why people consider this to be a very tedious process. To comprehend this, let's discuss a simple case of cryptography.

Let sender (S) send a message (called plaintext) to a particular receiver (R). 'S' converts his plaintext message to a secret form for transmission (which we may call the ciphertext) with the aid of a cryptographic algorithm (CA) and some defined key (K) parameters. 'CA' is the encryption method used here.

The whole procedure assumes some prior method of communication, as 'R' needs to know the parameters. The headache of the crypt analyst is that he needs to decipher the plaintext from the ciphertext without knowing the key parameters.

As we discussed, one of the simplest (and one of the oldest, too!) methods for encryption is the Caesar cipher method. Here, if a character in a particular place of the word is the  $N^{\text{th}}$  letter of the alphabet series, it is replaced by the  $(N + K)^{\text{th}}$  letter in the series, where K is the parameter-an integer (Caesar used  $K = 3$ !).

```
CAESAR(CA,N,K)
for_all_characters
character (N+K)? character(N)
```

You can add more statements to fix bugs (say, if you're using English, you can specify what to do if  $(N+K)$  exceeds 26).

Well, as said before, this method is very simple. Therefore, it's no big deal for the crypt analyst to crack the encrypted data. Things will become more complex if we use a general table to define the substitution and then use the same for the process. But here, too, our villain can try some tricks. He may choose the first character arbitrarily, say E (as E is the most frequent letter in English text). He may also choose not to go for certain diagrams such as QJ (as they never occur together in English).

You can develop the method further by using multiple look-up tables. Then, you will come across many interesting cases like the one when the key is as long as the plaintext ('one-time pad' case) and so on. It should be noted that if the message and key are encoded in binary, a more common scheme for position-by-position encryption is to use the "exclusive-or" function to encrypt the plaintext-"exclusive-or" it (bit by bit) with the key.

## Geometric algorithms

This methodology can be adopted to solve complex problems that are inherently geometric. It can be applied to solve problems concerning physical objects ranging from large buildings (design) and automobiles, to very large-scale integrated circuits (ICs).

But, you will soon see that even the most elementary operations (even on points) are computationally challenging. The interesting aspect is that some of these problems can readily be solved just by looking at them (and some others by applying the concepts in graph theory). If we resort to computational methods, we may have to go in for non-trivial methodologies.

This branch is relatively new and many fundamental algorithms are still being developed. Hence you can consider this as a potentially challenging and promising realm.

In this introductory piece, we'll restrict ourselves to the two-dimensional space. If you are able to properly define any point, then we can easily manage to include complex geometrical objects, say a line (as it is a pair of points connected by a straight line segment) or a polygon (defined by a set of points-array).

We can represent them by:

```
type point = record x,y: integer end;
line = record pl, p2: point end;
```

It is quite easy to work with pictures compared to numbers, especially when it comes to developing a new design (algorithm) pattern. It is also very helpful while debugging the code.

Let's see a recursive program that will enable us to 'draw' a line by drawing the endpoints.

```
procedure draw(l: line);
variable Δx, Δy: integer;
p: point; 10,11: line;
begin
dot(l.pl.x,l.pl.y); dot(l.p2.x,l.p2.y);
Δx:=l.p2.x-l.pl.x; Δy:=l.p2.y-l.pl.y;
if (abs(Δx)>1) or (abs(Δy)>1) then
begin
p.x:=l.pl.x+Δx div 2; p.y:=l.pl.y+Δy div 2;
l.pl:=l.pl; l.p2:=p; draw(l0);
l2.pl:=p; l2.p2:=l.p2; draw(l1);
end;
end;
```

You can see that there is a division of the space into two parts, joined by using line segments. You may stumble upon many algorithms where we will be converting geometric objects to points in a specific way. We can group them under the term 'scan-conversion algorithms'. To get a clear picture, you may write the pseudo code to check whether two lines are intersecting. (Hint: check for a common point.)

If you can't straight away do it, try this function to compute these lines and check whether they meet our condition:

```
function same_point(l: line; pl,p2: point): integer;
variable Δx, Δy, Δx1, Δx2, Δy1, Δy2: integer;
begin
Δx:=l.p2.x-l.pl.x; Δy:=l.p2.y-l.pl.y;
Δx1:=pl.x-l.pl.x; Δy1:=pl.y-l.pl.y;
Δx2:=p2.x-l.p2.x; Δy2:=p2.y-l.p2.y;
same_point:=(Δx*Δy1-Δy*Δx1)*(Δx*Δy2-Δy*Δx2)
end;
```

If the quantity  $(\Delta x, \Delta y_1 - \Delta y, \Delta x_1)$  is non-zero, we can say that  $pl$  is not on the line.

## A problem for beginners

Here we are not trying to address a real problem! We will look at how to produce graphical output with the help of libraries. You might have drawn 'pictures' in BASIC while at school, but this is not that method. In fact, our intentions are different.

Let's define our problem: We need to draw a sphere with the help of a few straight lines.

We can use HoloDraw (see the resource links for more information) as the library for drawing the sphere and we will do the codes in Shell. We start by 'flattening' the sphere to a flat rectangular map.

As it is a sphere, we will meddle with the changes in terms of 'degrees'. We also need an input file for processing by the HoloDraw. (Before you proceed, download a copy of HoloDraw and untar it into a local directory. Also make sure that you have Perl installed.)

The input file, *sphere.draw*, will be quite akin to the following:

```
color=0 1 0
draw a line around the sphere's equator
line: 0 0 1000, 360 0 1000
line: 0 45 1000, 360 45 1000
line: 0 -45 1000, 360 -45 1000

color=0 0 1
line: 0 90 1000, 0 -90 1000
line: 180 90 1000, 180 -90 1000

line: 30 90 1000, 30 -90 1000
line: 60 90 1000, 60 -90 1000
line: 90 90 1000, 90 -90 1000
line: 120 90 1000, 120 -90 1000
line: 150 90 1000, 150 -90 1000
line: 210 90 1000, 210 -90 1000
line: 240 90 1000, 240 -90 1000
line: 270 90 1000, 270 -90 1000
line: 300 90 1000, 300 -90 1000
```



## Finding (opting for) a strategy and the efficiency factor

While designing strategies it is important to consider their viability, effectiveness and efficiency. To comprehend the idea completely, consider a basic problem in quantum mechanics.

Schrödinger equation for the time-dependent wave function can be written as:

$$i\hbar \frac{\partial}{\partial t} |\Phi(t)\rangle = H |\Phi(t)\rangle$$

We can also write an expression for the thermal expectation value of an observable X as:

$$\langle X \rangle = \frac{\text{Tr} e^{-\beta H} X}{\text{Tr} e^{-\beta H}}$$

You can see that the above equation is modelled by a Hamiltonian H. Classically, it is quite easy to come out with a computational method to solve such equations (say by using Monte Carlo methods). But here the problem is that the objects (say operators or matrices) in QM do not necessarily commute.

Still, we can go for models defined by:

$$\langle X \rangle = \sum_{i=0}^{M-1} e^{-\beta E_i} \langle \Phi_i | X | \Phi_i \rangle / \sum_{i=0}^{M-1} e^{-\beta E_i}$$

A lattice of L sites filled with L/2 electrons with up spin, and L/2 electrons with down spin, is a physical model that easily fits into this. (Please Google the term 'Hubbard model' for more information about a better model.) But to find out what is really required to carry out these few steps, we need an order of magnitude for M. And by using approximation methods (like the Sterlings method) we can see that:

$$M \approx \frac{2^{2L+2}}{2\pi L}$$

This means that the quantity M increases exponentially with 2L (approximately). And if we allocate 8 bytes per floating point number, the amount of memory we need to store a single eigenvector will turn out to be:

$$\text{MEMORY} \approx \frac{2^{2L-25}}{2\pi L} \text{Gb}$$

So if I put L = 64, the memory required will be 1028 GB! This means that I need 1028 GB to study a quantum system of just 64 particles on 64 sites. If I submit a proposal with such high values, I am sure that no funding agency will accept this.

The only way I can do the computational task is to go for an algorithmic strategy that will reduce the amount of memory needed, at the expense of more CPU time. This is further considered in relation to 'clouds' and their effectiveness.

line: 330 90 1000, 330 -90 1000

Here the X and Y values (which you can identify from the codes directly) are in degrees around the sphere. And Z (or some axis reference) is the sphere's radius. As you can see, we have used different colours for east-west lines and north-south lines.

Now we will create our flat grid file from this, using the following shell code:

```
#!/bin/sh
/path_to_holodraw/drawwrl.pl < /location_of_input_file/sphere.draw >
flatgrid.wrl
```

But when we draw the sphere, we have to slice our long lines into small ones, so that our sphere will have a 'smooth' curve. We can do that by using the 'drawchop' and 'drawball' library files:

```
#!/bin/sh
/path_to_holodraw/drawchop.pl x=15+15 y=15+15 < /location_of_input_
file/sphere.draw |
/path_to_holodraw/drawball.pl |
/path_to_holodraw/drawwrl.pl > ballgrid.wrl
```

We can create the VRML (Virtual Reality Modelling Language) using the 'drawwrl' file:

```
#VRML V2.0 utf8
draw a line around the sphere's equator
Shape {
 appearance Appearance {
 material Material {
 emissiveColor 0 1 0
 transparency 0
 }
 }
 geometry IndexedLineSet {
 coord Coordinate {
 point [
 0 0 1000,
 500 0 866.025403784439,
 866.025403784439 0 500,
 1000 0 6.12303176911189e-14,
 866.025403784439 0 -500,
 500 0 -866.025403784439,
 1.22460635382238e-13 0 -1000,
 -500 0 -866.025403784439,
 -866.025403784438 0 -500,
 -1000 0 -1.83690953073357e-13,
 -866.025403784439 0 500,
 -500 0 866.025403784438,
 -2.44921270764475e-13 0 1000
]
 }
 coordIndex [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
 }
}
```

## Some 'tree' facts

In the last column, we discussed the use of trees. I will now list some of their properties that you can employ while designing the strategy:

- There will only be one node that connects two nodes in a tree
- If a tree has N nodes, there will be N-1 edges
- For any binary tree with N internal nodes, there are N+1 external nodes
- The height of a given full binary tree with N internal nodes is about  $\log N / \log 2$

```

}
}

Shape {
 appearance Appearance {
 material Material {
 emissiveColor 0 1 0
 transparency 0
 }
 }
 geometry IndexedLineSet {
 coord Coordinate {
 point [
 0 707.106781186547 707.106781186548,
 353.553390593274 707.106781186547 612.372435695795,
 612.372435695795 707.106781186547 353.553390593274,
 707.106781186548 707.106781186547 4.32963728535968e-14,
 612.372435695795 707.106781186547 -353.553390593274,
 353.553390593274 707.106781186547 -612.372435695795,
 8.65927457071935e-14 707.106781186547 -707.106781186548,
 -353.553390593274 707.106781186547 -612.372435695795,
 -612.372435695794 707.106781186547 -353.553390593274,
 -707.106781186548 707.106781186547 -1.2988911856079e-13,
 -612.372435695795 707.106781186547 353.553390593274,
 -353.553390593274 707.106781186547 612.372435695794,
 -1.73185491414387e-13 707.106781186547 707.106781186548
]
 }
 coordIndex [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]
 }
}
.....
.....

```

This way the code goes on. (The complete code of flatgid.wrl is available at [aasisvinayak.com/new\\_zone/forum.php?do=viewtopic&cat=2&topic=1](http://aasisvinayak.com/new_zone/forum.php?do=viewtopic&cat=2&topic=1))

We can generalize it as :

```

Shape {
 appearance Appearance {
 material Material {
 emissiveColor x x x

```

## Some evolutionary concepts: In a nutshell

Evolutionary algorithms themselves form another major branch. We will confine ourselves to some basic ideas, problem definitions and generalisations (definitions).

General single-objective optimisation problem: This is defined as minimising (or maximising)  $f(x)$  subject to  $g_i(x) \leq 0, i = \{1, \dots, m\}$ , and  $h_j(x) = 0, j = \{1, \dots, p\} x \in \Omega$ . A solution minimises (or maximises) the scalar  $f(x)$  where  $x$  is a  $n$ -dimensional decision variable vector  $x = (x_1, \dots, x_n)$  from some universe  $\Omega$ .

Single-objective global minimum optimisation: Given a function  $f : \Omega \subseteq R^n \rightarrow R, \Omega = \emptyset$ , for  $x \in \Omega$  the value  $f^* f(x^*) > -\infty$  is called a global minimum if and only if

$$\forall x \in \Omega : f(x^*) \leq f(x)$$

$x^*$  is by definition the global minimum solution,  $f$  is the objective function, and the set  $\Omega$  is the feasible region of  $x$ .

### Useful facts:

- The purpose of finding the global minimum solution(s) is called the global optimisation problem for a single-objective problem.
- Evolutionary multi-objective optimisation (EMO) refers to the use of evolutionary algorithms of any sort (like genetic algorithms, evolution strategies, evolutionary programming or genetic programming) to solve multi-objective optimisation problems.
- Other meta-heuristics that are being used to solve multi-objective optimisation problems include particle swarm optimisation, artificial immune systems and cultural algorithms.
- Differential evolution, ant colony, tabu search, scatter search, and memetic algorithms are other key ideas in the realm.

### Key ideas:

- You must see that non-dominated points are preserved in objective space, and the associated solution points in the decision space.
- The design should be such that it should continue to allow algorithmic progress towards the Pareto Front in the objective function space.
- Maintain the diversity of points on Pareto/phenotype front (space) or of Pareto optimal solutions on decision/genotype space.
- Provide the decision maker (DM) sufficient but limited number of Pareto points for the selection (which results in decision variable values).

Please let me know if you wish to discuss these ideas more in depth.

```

transparency x
}
}
geometry IndexedLineSet {
 coord Coordinate {
 point [

```

```

xyz
}
}
coordIndex [x,y,z]
}
}

```

...where x,y,z are local variables with respect to each reference point. And footer lines will be akin to:

```

#HISTORY# /home/aasisvinayak/Documents/Desktop/holodraw0.37/
drawchop.pl x=30+30 y=30+30
#HISTORY# /home/aasisvinayak/Documents/Desktop/holodraw0.37/
drawball.pl

```

```

NavigationInfo {
type ["EXAMINE", "FLY", "WALK", "ANY"]
speed 1.0
}
#HISTORY# /home/aasisvinayak/Documents/Desktop/holodraw0.37/
drawwrl.pl

```


This keeps track of the functions we employed.

[Initially, I thought of putting the generated images here, but later I felt that it was better to put the code itself because once you have the copy of the 'drawwrl' Perl source file, you can use it to analyse our input and the

corresponding output.]

We have seen that with the help of libraries, we can generate complex codes quite easily. So you can employ such functions, libraries and black-boxes when you write the algorithms.

If you are able to achieve this, then you can straight away try geometrical algorithms.

Having completed a good portion of our new segment, we can discuss the ideas you suggested. But I think it is too late to discuss notations (and advanced ideas in numerical computation) today. So wait for the forthcoming issues, in which we will address them. **END** 

### Resources

- <http://simkin.asu.edu/holodraw/download.html>
- <http://www.perl.com/>
- [http://www.dmoz.org/Computers/Software/Internet/Clients/VRML/Browser\\_Plugins/](http://www.dmoz.org/Computers/Software/Internet/Clients/VRML/Browser_Plugins/)
- [http://www.web3d.org/x3d/vrml/tools/viewers\\_and\\_browsers/](http://www.web3d.org/x3d/vrml/tools/viewers_and_browsers/)

### By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at [www.aasisvinayak.com](http://www.aasisvinayak.com)



Calling...

# Share your knowledge

## Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of [www.openITis.com](http://www.openITis.com) (aka [linuxforu.com](http://linuxforu.com)) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

[lfyedit@efyindia.com](mailto:lfyedit@efyindia.com)

### Few topics that top our list:

- ◆ Tips 'n' Tricks for software developets or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools

# Winners

## of Diwali Dhamaka



### 1<sup>st</sup> Prize Plasma TV



Mrs Poonam Kapoor, Director, EFY Group,  
presenting the Plasma TV to

**Ms Deepalaxmi**, the **First Prize Winner**



### 2<sup>nd</sup> Prize MP4 Player

1. Mr Vilas Katke, Mumbai
2. Mr Earnest Selva Paul, Secunderabad
3. Mr Rishi Kumar, Delhi
4. Mr Vijay Kumar, Bangalore
5. Sarada Instt of Tech. & Science, Khammam

6. Great Lakes Instt. of Management, Chennai
7. Mr Rahul Singh Kotesa, Goa
8. Collage of Engineering, Maharashtra
9. Bharathiyar Centenary Memori, Tamil Nadu
10. Ms S Padhmalakshmi, Chennai



### 3<sup>rd</sup> Prize Travel Bag

- |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ol style="list-style-type: none"> <li>1. Mr Viraj Patel, Mumbai</li> <li>2. Ms Shruti Verma, New Delhi</li> <li>3. Ranganathan Engineering, Coimbatore</li> <li>4. SRA Systems Ltd, Chennai</li> <li>5. Mr Sri Balaji M, Bangalore</li> <li>6. Mr Ravinder Reddy Tumu, Bodhan</li> <li>7. Mr Anuraj Anand, Pathanamthitta</li> <li>8. Akashganga AME India Pvt Ltd, Chennai</li> <li>9. Mr V Sundaresan, Chennai</li> <li>10. Mr Darshan Kumar, Punjab</li> <li>11. Mr Om Prakash Kanoongo, Mumbai</li> <li>12. Mr Jaisingh Varma, Mumbai</li> <li>13. Mr Niranjan Sahoo, Visakhapatnam</li> <li>14. Mr K Parthiban, Tiruchirappalli</li> </ol> | <ol style="list-style-type: none"> <li>16. Mr N M Irshath, Kavalpatnam</li> <li>17. Mr Vismay Buche, Hyderabad</li> <li>18. VVM's Shree Damodar College, Goa</li> <li>19. Mr P R Mantry, A P</li> <li>20. Dr Bhagwati Prasad, Span, Mumbai</li> <li>21. Mr Sivaji Mopidevi, EATON, Pune</li> <li>22. Mr Kishor Narkhede, Secunderabad</li> <li>23. Industrial Training Instt, Pune</li> <li>24. St Mira's College For Girls, Pune</li> <li>25. Dr P K P Mahamood, Kerala</li> <li>26. Mr Steve Antony Sequeira, Karnataka</li> <li>27. Indira Shiva Rao Polytechnic, Karnataka</li> <li>28. V N Krishnaswamy Naidu, Coimbatore</li> <li>29. Mr Narendra K Sangame, Karnataka</li> <li>30. Mr A Sudarshan, Virudhunagar</li> <li>31. Mr Ramkumar R, Coimbatore</li> <li>32. Government Industrial, Chittoor</li> <li>33. Mr Bijju Kumar J, Kerala</li> <li>34. Mr T Vetivel, Tamil Nadu</li> <li>35. Fr Nijo, Principal, Kerala</li> </ol> | <ol style="list-style-type: none"> <li>36. Perfect Communications, Ludhiana</li> <li>37. Rajeev Electronics Pvt Ltd, New Delhi</li> <li>38. The Vazir Sultan College of Engg., Khammam</li> <li>39. Mr R K Maharana, Orissa</li> <li>40. Mr Sajith Kumar V R, Bangalore</li> <li>41. Mr Pundalik Sutar, Mumbai</li> <li>42. Siemens Enterprises Communications Pvt Ltd, New Mumbai</li> <li>43. Mr Sunil R, Kerala</li> <li>44. Safa College Of Engg &amp; Tech, Andhra Pradesh</li> <li>45. Mr B N V Prasad, Andhra Pradesh</li> <li>46. Mr Venkatesh V, Chitradurga</li> <li>47. Mr B L Desai, Karnataka</li> <li>48. Mr Krishna Prasad Y Bhat, Karnataka</li> <li>49. Mr Biswanath Das, Kolkata</li> <li>50. Sivanandha Mills Ltd, Coimbatore</li> <li>51. Shriram Instt of Engg Technology, Maharashtra</li> <li>52. Ms Lalita Bhardwaj, Sonapat</li> <li>53. City Power Conversion, Secunderabad</li> <li>54. Mr Syan Kumar R, Cochin</li> <li>55. Mr Vinod Kumar P P, Kerala</li> <li>56. Mr Pankaj Bhagat, Hyderabad</li> <li>57. Mr Manoj Rakhyani, Madhya Pradesh</li> <li>58. VLB Janakiammal Polytechnic College, Tamil Nadu</li> </ol> | <ol style="list-style-type: none"> <li>59. Mr A.M. Panduranga, Karnataka</li> <li>60. Mr Manuvel Nadar M, Kerala</li> <li>61. Mr Raghavendra C, Karnataka</li> <li>62. Mr Bharat H Karani, Mumbai</li> <li>63. Mr V R Aneesh, Kerala</li> <li>64. Mr K Sathianarayanan, Chennai</li> <li>65. Sri Sankara Arts &amp; Science College, Tamil Nadu</li> <li>66. Mr Dinesh Maliya, Rajasthan</li> <li>67. Mr Labh Singh, Bathinda</li> <li>68. Info Instt of Engg, Coimbatore</li> <li>69. J R Communications &amp; Power Controls, Trichy</li> <li>70. Al-Madina College of Computer Science, Andhra Pradesh</li> <li>71. Ms Vijayalakshmi K, Bangalore</li> <li>72. Mr Deependra Kumar Rajput, Bangalore</li> <li>73. Mr Arvind Kumar, Kaithal</li> <li>74. Mr Ramakrishna V, Bangalore</li> <li>75. Centre for Environment Education, Ahmedabad</li> <li>76. Dhanalakshmi Srinivasan Engineering College, Tamil Nadu</li> <li>77. Sandur Polytechnic, Karnataka</li> <li>78. Mr Kirit P Budh, Gujarat</li> <li>79. Mr Sadiq Hussain, Dibrugarh</li> <li>80. Mr Kirit P Budh, Gujarat</li> </ol> | <ol style="list-style-type: none"> <li>81. Vimala College, Trichur</li> <li>82. Rotary Midtown Library, Rajkot</li> <li>83. Ms Snehal Joshi, Gujarat</li> <li>84. Mr Komel Bhojani, Pune</li> <li>85. The ICFAI Institute of Science &amp; Technology, Jaipur</li> <li>86. Mr Sachin G. Gune, Pandharpur</li> <li>87. Mr Z Jatin Shah, Mumbai</li> <li>88. R P Gogate College of Arts &amp; Science, Maharashtra</li> <li>89. St Edmund's College, Meghalaya</li> <li>90. Amity Electronics Corporation, Mathura</li> <li>91. Mr Bhaskar N Chhibber, Pune</li> <li>92. Sangamner Nagarpalika Arts, Maharashtra</li> <li>93. Mr Charan Jit Singh, New Delhi</li> <li>94. Mr Sumanta Sarkar, Gwalior</li> <li>95. Mr Sanjay V M, Bangalore</li> <li>96. Sree Narayana College of Technology, Kerala</li> <li>97. Mr Satish Kumar Bidwaik, Mumbai</li> <li>98. Mr Suresh S, Kerala</li> <li>99. Sri Venkata Ramaswamy, Karnataka</li> <li>100. Mr Tufan Sharma, Maharashtra</li> </ol> |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|





**PAY  
FOR  
ONE**

---

**GET  
TWO**

*Celebrating 41<sup>st</sup> Year*

**YOU CAN CHOOSE ANY OF THE FOLLOWING OPTIONS**

**Pay for 6 months  
Get for 1 year!**

**Pay for 1.5 years  
Get for 3 years!**

**Pay for 2.5 years  
Get for 5 years!**

### MEGA SUBSCRIPTION OFFER

| MAGAZINE                         | <input checked="" type="checkbox"/> Mark Required | Unit Price (Rs.) | Pay for 6 months<br>Get for 1 Year     | Pay for 1.5 Years<br>Get for 3 Years    | Pay for 2.5 Years<br>Get for 5 Years    |
|----------------------------------|---------------------------------------------------|------------------|----------------------------------------|-----------------------------------------|-----------------------------------------|
| Electronics For You (Without CD) | <input type="checkbox"/>                          | 40/-             | <input type="checkbox"/> Pay Rs. 250/- | <input type="checkbox"/> Pay Rs. 730/-  | <input type="checkbox"/> Pay Rs. 1210/- |
| Electronics For You (with CD)    | <input type="checkbox"/>                          | 80/-             | <input type="checkbox"/> Pay Rs. 380/- | <input type="checkbox"/> Pay Rs. 1080/- | <input type="checkbox"/> Pay Rs. 1800/- |
| Electronics Bazaar               | <input type="checkbox"/>                          | 50/-             | <input type="checkbox"/> Pay Rs. 310/- | <input type="checkbox"/> Pay Rs. 910/-  | <input type="checkbox"/> Pay Rs. 1510/- |
| Linux For You (with CD & DVD)    | <input type="checkbox"/>                          | 100/-            | <input type="checkbox"/> Pay Rs. 610/- | <input type="checkbox"/> Pay Rs. 1810/- | <input type="checkbox"/> Pay Rs. 3010/- |
| BenefIT                          | <input type="checkbox"/>                          | 60/-             | <input type="checkbox"/> Pay Rs. 300/- | <input type="checkbox"/> Pay Rs. 900/-  | <input type="checkbox"/> Pay Rs. 1500/- |
| Facts For You                    | <input type="checkbox"/>                          | 100/-            | <input type="checkbox"/> Pay Rs. 800/- | <input type="checkbox"/> Pay Rs. 1800/- | <input type="checkbox"/> Pay Rs. 3000/- |

• These rates are applicable for new subscribers as well as renewal by existing subscribers. • The rates are valid for subscribers within India only. • Please allow 4-8 weeks for processing of your subscription. • Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

**Hurry!**

Offer Valid till 31<sup>st</sup> March 2009



A magazine for the electronics fraternity... professionals, businessmen and hobbyists



India's first Electronics Sourcing magazine



Asia's first magazine on Linux and Open Source



An IT Magazine for Business Decision Makers



A monthly on business & economic affairs with regular sections on market surveys

# Subscribe Now!

## ORDER FORM

Name \_\_\_\_\_ Company Name \_\_\_\_\_ Designation \_\_\_\_\_

Mailing Address \_\_\_\_\_ Dist. \_\_\_\_\_

City \_\_\_\_\_ Pin Code \_\_\_\_\_ State \_\_\_\_\_ Phone \_\_\_\_\_

Email \_\_\_\_\_ Subscription No. (for existing subscribers only) \_\_\_\_\_

I would like to subscribe to the above ☒ printed magazine(s), for the period \_\_\_\_\_, Please find enclosed a sum of Rs \_\_\_\_\_ by DD/MQ/crossed cheque\*bearing the No. \_\_\_\_\_ dt. \_\_\_\_\_ in favour of EFY Enterprises Pvt Ltd, payable at Delhi. (\*\*Please add Rs 50 on non-metre cheque)

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. \_\_\_\_\_

against my credit card No. \_\_\_\_\_

Date of Birth \_\_\_\_/\_\_\_\_/\_\_\_\_ (dd/mm/yy) Card Expiry Date \_\_\_\_/\_\_\_\_/\_\_\_\_ (mm/yyyy)



Signature (as on the card)



Send this filled-in form or its photocopy to :

EFY Enterprises Pvt Ltd: D-87/1, Oldia Industrial Area, Phase I, New Delhi 110 020

Ph: 011-26810601-03; Fax: 011-26817583, E-mail: info@efyindia.com website: www.efyindia.com



# LINUX JOBS

## Post: QA Lead

Company: Intelligaia Technologies Pvt. Ltd.  
Profile: Hands on experience on Linux/ Unix/ Windows/ Web applications/ Distributed applications, .net, C, C++ Programming, Java Coding, GUI testing, SDK level Testing with an experience in black Box testing is desired.  
Exp: 5-10  
Location: Chandigarh  
Email: careers1@intelligaia.com

## Post: Oracle DBA

Company: Spectraforce Technologies (India) Pvt. Ltd.  
Profile: Candidate should have 4 years of experience. Should have worked on RMAN besides Windows environment & should be having an experience on Unix, Linux.  
Exp: 4-8  
Location: Bangalore  
Email: smanga@spectraforce.com

## Post: Project Leader- Voyager

Company: Atlanta India Pvt. Ltd.  
Profile: Must have hands on in C/C++, with work experience in data-structures plus good knowledge of one of the HDL languages (Verilog, VHDL). Should be B.E./ B.Tech/ M. Tech. in Computer Science/ Electrical Electronics Engineering from reputed Universities.  
Exp: 5-6  
Location: Noida  
Email: career@noida.atlanta.com

## Post: Network and Infrastructure Manager

Company: Eastern Software Systems Pvt. Ltd.  
Profile: Selected candidate will be responsible for leading and managing global network systems and designing and building of CISCO infrastructures (LAN/ WAN) and in-depth understanding of routing protocols. Proficiency in Cisco technologies and IP networks is desired.  
Exp: 7-10  
Location: Kenya  
Email: ess-delhi@essindia.co.in

## Post: Japanese Language Known DBA's required

Company: Adea Technologies Pvt. Ltd.  
Profile: Candidate should have 4-9 yrs of prior working experience.  
Exp: 4-9  
Location: Bengaluru  
Email: Anita.N@adea.com

## Post: Production Support DBA, Oracle, SQL Server

Company: Diaspark Inc.  
Profile: Bachelors/ Master's degree in Computer Science or related field with at least 5 plus years of experience is desired. Minimum 5 years of experience as Oracle/ SQL Server DBA is required.  
Exp: 5-10  
Location: United States (U.S.)  
Email: nagesh.tiwari@diaspark.com

## Post: Sr. Oracle DBA

Company: Akshay Software Technologies Ltd.  
Profile: ORACLE DBA technical skills in AIX and Solaris Unix, Linux and Windows Platform in production environments in Oracle 9i, 9g and 10g versions is required.  
Exp: 6-7  
Location: Delhi  
Email: hrd@akshay.com

## Post: EMC Storage Management Systems Administrator

Company: Source One Management Services Pvt. Ltd.  
Profile: Bachelor's Degree or Diploma in Computer Science in Electronics or relevant experience in EMC Storage is desired. 3+ years of experience is must.  
Exp: 4-9  
Location: Noida  
Email: amit.mishra@sourceone.co.in

## Post: Sr. Unix Admin

Company: SAIC India Pvt Ltd  
Profile: 7+ years experience in Unix/ Linux system administration or 4 years relevant experience and a B.S. in computer science field is desired.  
Exp: 6-10  
Location: Bengaluru  
Email: pappus@saic.com

## Post: Database Administrator and System Admin

Company: Newt Global Consulting, LLC  
Profile: Experience with Tier-1 & Tier-2 companies in enterprise environment ranging from 4-10 yrs is desired.  
Exp: 4-9  
Location: United States (U.S.)  
Email: k1bjobs2@newtglobal.com

## Post: Sr. Application Engineer

Company: Amazon Development Centre Pvt. Ltd.  
Profile: Ideal candidate should be B.S. in Computer Science or a related field plus 5-7 yrs of overall development/ technical support experience. Detailed knowledge of the UNIX/ Linux operating system and tools is also desired.  
Exp: 4-8, Location: Chennai  
Email: withyan@amazon.com

## Post: Storage Admin

Company: Omnitel Infotolutions Ltd.  
Profile: A BS/ MS in Computer Science, Information Systems, Engineering having at least 6 years track record in consulting/ IT management. Knowledge of operating Systems preferably in Unix, HP-UX, Solaris, AIX, Windows Server, Linux is desired.  
Exp: 7-12  
Location: Mumbai  
Email: alen@omnitelindia.com

## Post: Team Lead- PHP

Company: Iridium Interactive Ltd.  
Profile: Should possess at least 5-7 years of relevant experience. Will be responsible for preparation of use cases, understanding the client requirement & studying on the technical feasibility.  
Exp: 5-7  
Location: Hyderabad  
Email: hr@iridiuminteractive.com

## Post: AIX Admin- Technical Lead

Company: TESCO Hindustan Service Centre  
Profile: Good experience in system implementation and administration in AIX is desired having more than 4 years experience in cluster implementation and administration (HACMP). Knowledge of other flavors of Unix like Linux (Redhat and SUSE) and HP-UX administration is also desired.  
Exp: 8-12  
Location: Bangalore  
Email: nagesh.kasaragod@in.tesco.com

## Post: Systems Engineer - Linux Servers

Company: Netmagic Solutions Pvt. Ltd.  
Profile: Experience with multiple unix systems - Sun Solaris, HP-UX, AIX, Red Hat, and Debian is desired. Fine tuning of Applications- Websphere, Apache, JBoss, MySQL etc. is also desired.  
Exp: 5-7  
Location: Mumbai  
Email: juhi.shah@netmagicsolutions.com

## Post: Unix (Linux, Solaris, AIX, HP/UX) Systems Admins

Company: Software Pundits  
Profile: Seeking highly motivated unix System Administrators (UNIX, AIX, HP-UX or Solaris) who have demonstrated experience of 8+ years covering all aspects of the UNIX infrastructure environment.  
Exp: 5-8  
Location: United States (U.S.)  
Email: globaljobs@pundits.com

## Post: Systems Engineer - Linux

Company: IVY Comptech Pvt. Ltd.  
Profile: job requires a Linux administrator with 4+ years of experience, with minimum 2+ on Linux/UNIX environment. Extensive knowledge of basics Unix concepts, file systems, permissions, ACLs, kernel configuration is desired.  
Exp: 4-9  
Location: Hyderabad  
Email: neerajas@ivycomptech.com

## Post: Sr. Oracle Applications- DBA

Company: Blink Consulting Pvt. Ltd.  
Profile: Should have Oracle Apps DBA experience of minimum 5 years (11.5.9/11.5.10.2) plus minimum 2 years of Lead experience. Must have worked on all flavors of operating system like HP/UX, IBM, Sun Solaris, Linux and Windows.  
Exp: 5-9  
Location: Pune  
Email: hr@blinkconsulting.com

## Post: Technical Manager

Company: WOW Global India Pvt. Ltd.  
Profile: An experience in citrix, vmware, veritas is desired. Certification is must in any windows, linux, unix, solaris. Japanese speaking is must.  
Exp: 6-10  
Location: Bengaluru  
Email: resourcing@wowglobal.com



**Get the Best Paying Jobs @ naukri.com**  
**2,00,000 Jobs to choose from!**

35,000 Recruiters | Get Jobs by email | Block Your Current Employers

Powered by  
**naukri.com**  
India's No.1 Job Site

For sample resumes/cover letters/Interview tips visit [www.naukri.com](http://www.naukri.com) or email us at [hr@naukri.com](mailto:hr@naukri.com)

# FOSS Yellow Pages

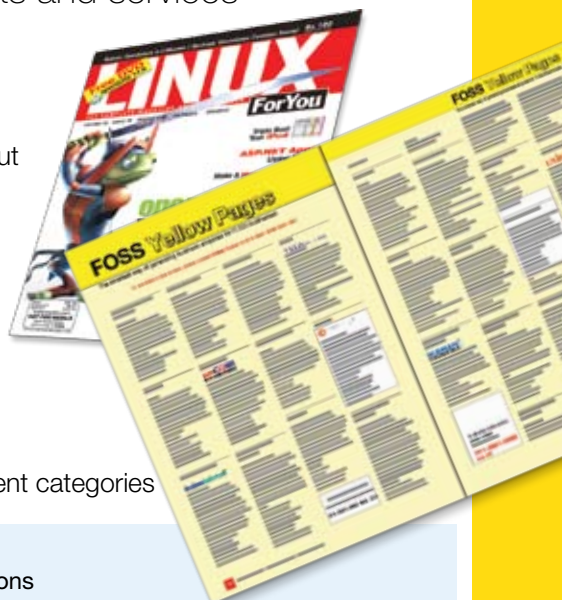
The best place for you to buy and sell FOSS products and services

## HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

## FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



## TARIFF

### Category Listing

|                           |          |
|---------------------------|----------|
| ONE Category .....        | Rs 2,000 |
| TWO Categories.....       | Rs 3,500 |
| THREE Categories.....     | Rs 4,750 |
| ADDITIONAL Category ..... | Rs 1,000 |

### Value-add Options

|                                         |          |
|-----------------------------------------|----------|
| LOGO-plus-Entry.....                    | Rs 500   |
| Highlight Entry (white background)..... | Rs 1,000 |
| Per EXTRA word (beyond 30 words).....   | Rs 50    |

## KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

## TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

✂ Tear & Send

## ORDER FORM

Tear & Send ✂

Organisation Name (70 characters): \_\_\_\_\_

Description (30 words): \_\_\_\_\_

Email: \_\_\_\_\_ Website: \_\_\_\_\_

STD Code: \_\_\_\_\_ Phone: \_\_\_\_\_ Mobile: \_\_\_\_\_

Address (will not be published): \_\_\_\_\_

City/Town: \_\_\_\_\_ Pin-code: \_\_\_\_\_

## CATEGORIES

- |                                                             |                                                      |                                                     |
|-------------------------------------------------------------|------------------------------------------------------|-----------------------------------------------------|
| <input type="checkbox"/> CONSULTANTS                        | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING  | <input type="checkbox"/> SOFTWARE DEVELOPMENT       |
| <input type="checkbox"/> CONSULTANT (FIRM)                  | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS                 | <input type="checkbox"/> LINUX-BASED WEB-HOSTING     | <input type="checkbox"/> TRAINING FOR CORPORATE     |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS            | <input type="checkbox"/> THIN CLIENT SOLUTIONS      |

Please find enclosed a sum of Rs. \_\_\_\_\_ by DD/ MO//crossed cheque\* bearing the No. \_\_\_\_\_ dt. \_\_\_\_\_ in favour of

**EFY Enterprises Pvt Ltd, payable at Delhi. (\*Please add Rs. 50 on non-metro cheque) towards the cost of \_\_\_\_\_ FOSS Yellow Pages advertisement(s)**

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. \_\_\_\_\_

against my credit card No.

C V V No. \_\_\_\_\_ (Mandatory)

Date of Birth \_\_\_\_ / \_\_\_\_ / \_\_\_\_ (dd/mm/yy) Card Expiry Date \_\_\_\_ / \_\_\_\_ (mm/yy)

Signature (as on the card)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020  
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

**To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'lore: 09986075717)**



# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

## Consultant (Firm)

### OS3 Infotech

•Silver Solutions Partner for Novell  
•High Availability Computing  
Solutions •End-to-end Open Source  
Solutions Provider •Certified Red  
Hat Training Partner •Corporate and  
Institutional Training

Navi Mumbai

Mobile: 09324113579

Email: info@os3infotech.com

Web: www.os3infotech.com

### Taashee Linux Services

100% Support on LINUX ,OSS  
& JBOSS related projects. We  
specialize in high-availability and  
high-performance clusters,remote  
and onsite system management,  
maintenance services,systems  
planning, Linux & JBOSS consulting  
& Support services.

Hyderabad

Mobile: 09392493753, Fax: 040-40131726

Email: manojkumar@taashee.com

Web: www.taashee.com

## Computer (UMPC) For Linux And Windows

### Comptek International

World's smallest computer comptek  
wibrain B1 UMPC with Linux,Touch  
Screen, 1 GB RAM 60GB, Wi-Fi,  
Webcam, upto 6 hour battery  
(opt.), USB Port, max 1600x1200  
resolution, screen 4.8", 7.5"x3.25"  
Size, weight 526 gm.

New Delhi

Mobile: 09968756177, Fax: 011-26187551

Email: comptekdelhi@compteki.com

Web: www.compteki.com

or www.compteki.in

## Education & Training

### Aptech Limited

IT, Multimedia and Animation  
Education and Training

Mumbai

Tel: 022-28272300, 66462300

Fax: 022-28272399

Email: customercare@aptech.ac.in

Web: www.aptech-education.com,

www.arena-multimedia.com

### IT-Campus : Academy of Information Technology

IT training and solution company with  
over 12 years of experience. - RHCE

•Software Training •Hardware  
Training •Multimedia And Animation

•Web Designing •Financial  
Accounting

Kota (Raj.)

Tel: 0744-2503155, Mobile: 09828503155

Fax: 0744-2505105

Email: m\_trilok@yahoo.com

Web: www.doeacc4u.com

### Mahan Computer Services (I) Limited

Established in 1990, the organization  
is primarily engaged in Education  
and Training through its own &  
Franchise centres in the areas of IT  
Software, Hardware, Networking,  
Retail Management and English. The  
institute also provides customized  
training for corporates.

New Delhi

Tel: 011-25916832-33

Email: info@mahanindia.com

Web: www.mahanindia.com

## Enterprise Communication Solutions

### Aware Consultants

We specialize in building and  
managing Ubuntu/Debian  
Linux servers and provide good  
dependable system administration.  
We install and maintain in-house  
corporate servers. We also provide  
dedicated and shared hosting as well  
as reliable wireless/hybrid networking.

Bangalore

Tel: 080-26724324

Email: sales@aware.co.in

Web: www.aware.co.in

### ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty  
at the Indian Institute of Science,  
Bangalore and carry over eight  
decades of experience and  
fundamental knowledge in the field  
of DSP and Telecommunication.  
ESQUBE plays a dominant role in  
the creation of IP in the domain of  
Sensors, Signals and Systems.

Bangalore

Tel: 080-23517063

Email: info@esqube.com

Web: www.esqube.com

### Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open  
source Enterprise Groupware  
Mobile push, Email Syncing of  
Contacts/Calendar/Tasks with  
mobiles

•Mail Archival •Mail Auditing •Instant  
Messaging

New Delhi

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

## netCORE The Innovation Company

### Netcore Solutions Pvt Ltd

No.1 company for providing Linux  
Based Enterprise Mailing solution  
with around 1500+ Customer all over  
India. Key Solutions:

•Enterprise Mailing and Collaboration  
Solution •Hosted Email Security •Mail  
Archiving Solution •Push Mail on  
Mobile •Clustering Solution

Mumbai

Tel: 022-66628000

Mobile: 09322985222

Email: kalpit@netcore.co.in

Web: www.netcore.co.in



### Red Hat India Pvt Ltd

Red Hat is the world's leading  
open source solutions provider.  
Red Hat provides high-quality,  
affordable technology with its  
operating system platform, Red  
Hat Enterprise Linux, together with  
applications, management and  
Services Oriented Architecture (SOA)  
solutions, including JBoss Enterprise  
Middleware. Red Hat also offers  
support, training and consulting  
services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

## Hardware & Networking Institute

### Xenitis Technolab Pvt Ltd

Xenitis TechnoLab is the first of its  
kind, state-of-the-art infrastructure,  
Hardware, Networking and  
I.T Security training institution  
headquartered in Kolkata. TechnoLab

is the training division of Xenitis group  
of Companies. It is the proud owner  
of 'Aamar PC', the most popular  
Desktop brand of Eastern India.  
These ranges of PC's are sold in  
the west under the brand name of  
'Aamchi PC', in the north as 'Apna  
PC' and in the south as 'Namma PC'.

Kolkata

Tel: 033-22893280

Email: srinku@xenitisgroup.com

Web: www.technolabindia.com

## IT Infrastructure Solutions

### Absolut Info Systems Pvt Ltd

Open Source Solutions Provider.  
Red Hat Ready Business Partner.  
Mail Servers/Anti-spam/GUI  
interface/Encryption, Clustering &  
Load Balancing - SAP/Oracle/Web/  
Thin Clients, Network and Host  
Monitoring, Security Consulting,  
Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873839960

Email: sales@aisplglobal.com

Web: www.aisplglobal.com

### Advent Infotech Pvt Ltd

Advent has an experienced techno-  
marketing team with several years of  
experience in Networking & Telecom  
business, and is already making  
difference in market place. ADVENT  
qualifies more as Value Added  
Networking Solution Company, we  
offers much to customers than just  
Routers, Switches, VOIP, Network  
Management Software, Wireless  
Solutions, Media Conversion, etc.

New Delhi

Tel: 46760000, 09311166412

Fax: 011-46760050

Email: marketingsupport@

adventelectronics.com

Web: www.adventelectronics.com

### Asset Infotech Ltd

We are an IT solution and training  
company with an experience of 14  
years, we are ISO 9001: 2000. We  
are partners for RedHat, Microsoft,  
Oracle and all Major software  
companies. We expertise in legal  
software ans solutions.

Dehradun

Tel: 0135-2715965, Mobile: 09412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

To advertise in this section, please contact  
Somaiah (Bangalore) **09986075717**  
Dhiraj (Delhi) **09811206582**

**FOSS Yellow Pages**

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

**Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717**

## BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

## Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

## Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veritas, Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

**FOSS Yellow Pages**

## HBS System Pvt Ltd

System Integrators & Service Provider. Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symantec Partners of SUN for their new startup E-commerce initiative Solution Provider on REDHAT, SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amittal@hbsindia.com



## Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

## Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

## LDS Infotech Pvt Ltd

Is the authorised partner for RedHat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

## Pacer Automation Pvt Ltd

Pacer is leading providers of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, Vmware, Microsoft and Symantec. Our core expertise exists in, Consulting, building and Maintaining the Complete IT Infrastructure.

Bangalore

Tel: 080-42823000, Fax: 080-42823003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



## Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

## Srijan Technologies Pvt Ltd

Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

## TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

## Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of industry vehicles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tetra.in

Web: www.tetra.in

## Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

**FOSS Yellow Pages**

Want to register your organisation in  
FOSS Yellow Pages For **FREE\***

**Call:** Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717  
**or mail:** dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

*\*Offer for limited period.*

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

**Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717**

## Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

### Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

## Linux-Based Web-Hosting

### Manas Hosting

ManasHosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

### Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

## Linux Desktop

### Indserve Infotech Pvt Ltd

OpenLx Linux with Kalcuate (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

### New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

## Linux Experts

### Intaglio Solutions

We are the training and testing partners of RedHat and the first to conduct RHCS exam in delhi for the first time ever.

### New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

## Linux Vendor/Distributors

### GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

### Bangalore

Mobile: +91-9845009939, +91-9343861758

Email: sales@gte-india.com

Web: www.gte-india.com

### Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details

### Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

## Software Subscriptions

### Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.

### Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

## Software Development

### Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

### Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

## DeepRoot Linux Pvt Ltd

DeepRoot Linux is a seven year old GNU/Linux and Free Software company based in Bangalore. We develop Free Software products that are quick-to-deploy and easy-to-use.

### Bangalore

Tel: 080-40890000

Email: shivanand@deeproot.in

Web: www.deeproot.in



## InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

### Noida

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: http://opensource.infoaxon.com

## Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

### Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

## iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

### Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Want to register your organisation in  
**FOSS Yellow Pages**  
For **FREE**  
CALL: DHIRAJ (DELHI) 09811206582  
SOMAIAH (BANGALORE) 09986075717  
\*Offer for limited period.

## Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

### New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

## Software and Web Development

### Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

### Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

## CATEGORIES FOR FOSS YELLOW PAGES

- Consultants
- Consultant (Firm)
- Embedded Solutions
- Enterprise Communication Solutions
- High Performance Computing
- IT Infrastructure Solutions
- Linux-based Web-hosting
- Mobile Solutions
- Software Development
- Training for Professionals
- Training for Corporate
- Thin Client Solutions

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

**Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717**

## Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

### Patiala

Mobile: 91-9780531682

Email: [pardeep@mrsite.co.in](mailto:pardeep@mrsite.co.in)

Web: [www.mrsite.co.in](http://www.mrsite.co.in)

## Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

### New Delhi

Tel: 011-41648668, 66091565

Email: [ceo@salahsoftware.com](mailto:ceo@salahsoftware.com)

Web: [www.salahsoftware.com](http://www.salahsoftware.com)

## Thin Client Solutions

### Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing

Mobile: 09880715253

Email: [ranga@digitalwaves.in](mailto:ranga@digitalwaves.in)

Web: [www.digitalwaves.in](http://www.digitalwaves.in)

## Enjoy Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment.

Active Dealer Channel all over India.

### Gujarat

Tel.: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: [info@enjoyworld.com](mailto:info@enjoyworld.com)

Web: [www.enjoyworld.com](http://www.enjoyworld.com)

## Training for Corporate

### Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

### Ahmedabad

Tel: 079-27545455—66

Fax: 079-27545488

Email: [info@bascombridge.com](mailto:info@bascombridge.com)

Web: [www.bascombridge.com](http://www.bascombridge.com)

### Brainnet

#### Kolkata

Tel: 033-40076450

Email: [brainnet@brainware-india.com](mailto:brainnet@brainware-india.com)

Web: [www.brainware-india.com](http://www.brainware-india.com)

## Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

### Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586

Fax: 022-25706700

Email: [contact@cettm.mtnl.in](mailto:contact@cettm.mtnl.in)

Web: <http://cettm.mtnl.in/infra>

## Complete Open Source Solutions

RHCT, RHCE and RHCSS training.

### Hyderabad

Tel: 040-66773365, 9849742065

Email: [nayak.sujeet@gmail.com](mailto:nayak.sujeet@gmail.com)

Web: [www.cossindia.com](http://www.cossindia.com)

### ElectroMech

Redhat Linux and open source solution, RHCE, RHCSS training and exam center, Ahmedabad and Vadodara

### Ahmedabad

Tel: 079-40027898

Email: [electromech@electromech.info](mailto:electromech@electromech.info)

Web: [www.electromech.info](http://www.electromech.info)

## Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

### Cochin

Tel: 0484-2335324

Email: [enquiry@focuzinfotech.com](mailto:enquiry@focuzinfotech.com)

Web: [www.focuzinfotech.com](http://www.focuzinfotech.com)

## Gujarat Infotech Ltd

GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, Mysql

### Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: [info@gujaratinfotech.com](mailto:info@gujaratinfotech.com)

Web: [www.gujaratinfotech.com](http://www.gujaratinfotech.com)

## Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

### Chennai

Tel: 044-42171278, 9840880558

Email: [contactus@lynusacademy.com](mailto:contactus@lynusacademy.com)

Web: [www.lynusacademy.com](http://www.lynusacademy.com)

## Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

### Bangalore

Tel: 080-22428538, 26600839

Email: [info@linuxlearningcentre.com](mailto:info@linuxlearningcentre.com)

Web: [www.linuxlearningcentre.com](http://www.linuxlearningcentre.com)

## Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking,

while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

### Chennai

Tel: 044-45582525

Email: [info@mazenetsolution.com](mailto:info@mazenetsolution.com)

Web: [www.mazenetsolution.com](http://www.mazenetsolution.com)

## Netweb Technologies

Simplified and scalable storage solutions.

### Bangalore

Tel: 080-41146565, 32719516

Email: [info@netwebindia.com](mailto:info@netwebindia.com)

Web: [www.netwebindia.com](http://www.netwebindia.com)

## New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning. Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

### New Delhi

Tel: 011-43612400

Email: [info@nhindia.com](mailto:info@nhindia.com)

Web: [www.nhindia.com](http://www.nhindia.com)

## Network NUTS

India's only Networking Institute by Corporate Trainers. Providing



# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services



# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

**Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717**

Corporate and Open classes for RHCE / RHCSS training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

**New Delhi**

**Tel: 46526980-2**

**Mobile: 09310024503, 09312411592**

**Email: info@networknuts.net**

**Web: www.networknuts.net**

## STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

**New Delhi**

**Tel: 011-40560941-42, Mobile:**

**09873108801**

**Email: rakhi@stg.in**

**Web: www.stgonline.com**

**www.stgglobal.com**

## TNS Institute of Information Technology Pvt Ltd

Join RedHat training and get 100% job guarantee. World's most respected Linux certification. After RedHat training, you are ready to join as a Linux Administrator or Network Engineer.

**New Delhi**

**Tel: 011-3085100, Fax: 30851103**

**Email: nks@tiit.co.in**

**Web: www.tiit.co.in**

## Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

**Kolkata**

**Tel: 033-22833568, Mobile: 09433111110**

**Email: enquiry@webelinformatics.com**

**Web: www.webelinformatics.com**

To advertise in this section,  
please contact

**Somaiah (Bangalore)**

**09986075717**

**Dhiraj (Delhi)**

**09811206582**

## Training for Professionals

### Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

**Dehradun**

**Tel: 0135-2673712, Mobile: 09760099050**

**Web: www.agamtecindia.com**

### Amritha Institute of Computer Technology

Amritha Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.

**Mobile: 09393733174**

**Email: aict.hybd@amrita.ac.in**

**Web: www.amritahyd.org**

### Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.

**Bhopal**

**Tel: 0755-2661412, 2661559**

**Fax: 0755-4220022**

**Email: crisp@crispindia.com**

**Web: www.crispindia.com**

### Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

**New Delhi**

**Mobile: 09312506496**

**Email: info@cfosdr.com**

**Web: www.cfosdr.com**

### Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE,

RHCSS, CCNA, MCSE

**Kolkata**

**Tel: 033-25395508, Mobile: 09831705913**

**Email: info@cisconetinfo.com**

**Web: www.cisconetinfo.com**

### CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available..Leading Networking Institute in Marathwada

**Aurangabad**

**Tel: 0240-3299509, 6621775**

**Email: aurangabad@cmsinstitute.co.in**

**Web: www.cmsaurangabad.com**

### Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

**Bikaner**

**Tel: 0151-2202105, Mobile: 09928173269**

**Email: cmtech.bikaner@gmail.com,**

**kr.gupta.ashish@gmail.com**

### Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer Hardware & Networking

**Dehradun**

**Tel: 3208054, 09897168902**

**Email: thedishainstitute@gmail.com**

**Web: www.unisofttechnologies.com**

### EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux- Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

**Chandigarh**

**Tel: 0172-5067566-67, 2609849**

**Fax: 0172-2615465**

**Email: info@technoschool.net**

**Web: http://technoschool.net**

### GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications

**Jaipur**

**Tel: 0141-3213378**

**Email: franchise\_gt@gteducation.net**

**Web: www.gteducation.net**

### HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

**Bhopal**

**Tel: 0755-4094852**

**Email: bhopal@hclcdc.in**

**Web: www.hclcdc.in**

### IINZTRIX E Technologies Pvt Ltd

No. 1 Training provider in this region.

**meerut**

**Tel: 0121-4020111, 4020222**

**Mobile: 09927666664**

**Email: jai@iintrix.com**

**Web: www.iintrix.com**

### Indian Institute of Job Oriented Training Centre

**Ahmedabad**

**Tel: 079-40072244—2255—2266**

**Mobile: 09898749595**

**Email: info@iijit.net**

**Web: www.iijit.net**

### Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking  
•Software Engg. •Multimedia  
Training.

**Ahmedabad**

**Tel: 079-32516577, 26607739**

**Fax: 079-26607739**

**Email: contact@iantindia.com**

**Web: www.iantindia.com**

### IPCC

Bridging Gap with professionals.

**Lucknow**

**Tel: 0522-3919496**

**Email: ipccko@yahoo.co.in**

**Web: www.ipcc.co.in**

### IPSR Solutions Ltd

Earn RHCE / RHCSS certification, in Kerala along with a boating & free accommodation. IPSR conducted more than 2000 RHCE exams with 95-100% pass rate. Our faculty panel consists of 15 Red Hat Certified Engineers.

**Kochi, Kerala**

**Tel: +91 9447294635**

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

**Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717**

Email: [training@ipsrsolutions.com](mailto:training@ipsrsolutions.com)  
Web: [www.ipsr.org](http://www.ipsr.org)

## Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.  
**New Delhi**  
Mobile: 09910710143, Fax: 011-25886909  
Email: [info@koenig-solutions.com](mailto:info@koenig-solutions.com)  
Web: [www.koenig-solutions.com](http://www.koenig-solutions.com)

## NACS/CIT

We are Providing Training of LINUX to Professional & Cooperate.  
**Meerut**  
Tel: 0121-2420587, Mobile: 9997526668  
Email: [Info@nacsglobal.com](mailto:Info@nacsglobal.com)  
Web: [www.nacsglobal.com](http://www.nacsglobal.com)

## NACS Infosystems (P) Ltd

NACS is a organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.  
**Meerut**  
Tel: 0121-2767756, Fax: 0121-4006551  
Mobile: 09897796603  
Email: [info@nacsglobal.com](mailto:info@nacsglobal.com), [mohit@nacsglobal.com](mailto:mohit@nacsglobal.com)  
Web: [www.nacsglobal.com](http://www.nacsglobal.com)

## Netdiox Computing Systems

We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.  
**Bangalore**  
Tel: 080-26640708  
Mobile: 09740846885  
Email: [info@netdiox.com](mailto:info@netdiox.com)

## NetMax-Technologies

Training Partner of RedHat, Cisco  
**Chandigarh**  
Tel: 0172-2608351, 3916555  
Email: [mail.netmax@gmail.com](mailto:mail.netmax@gmail.com)  
Web: [www.netmaxtech.com](http://www.netmaxtech.com)

## Netxpri institute of Advance Networking

Netxpri Noida is a Leading organization to provide Open Source training on Red-Hat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.  
**Noida**  
Tel: 0120-4346847, Mobile: 09268829812  
Email: [anoop.vyas@netxpriindia.com](mailto:anoop.vyas@netxpriindia.com)  
Web: [www.netxpriindia.com](http://www.netxpriindia.com)

## Netzone Infotech Services Pvt Ltd

Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.  
**New Delhi**  
Tel: 011-46015674, Mobile: 9212114211  
Email: [info@netzoneindia.net](mailto:info@netzoneindia.net)

## Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.  
**Chennai**  
Tel: 044-2433 7355  
Email: [training@plexus.co.in](mailto:training@plexus.co.in)  
Web: [www.plexus.co.in](http://www.plexus.co.in)

## Professional Group of Education

RHCE & RHCSS Certifications  
**Jabalpur**  
Tel: 0761-4039376,  
Mobile: 09425152831  
Email: [naidu.vikas@gmail.com](mailto:naidu.vikas@gmail.com)

## Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.  
**Bangalore**  
Tel: 080-26639207, 26544135, 22440507  
Mobile: +91 9945 282834  
Email: [counsellors@qsoftindia.com](mailto:counsellors@qsoftindia.com)  
Web: [www.qsoftindia.com](http://www.qsoftindia.com)

## Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.  
**Chandigarh**  
Tel: 0172-5086829  
Email: [stn2001@rediffmail.com](mailto:stn2001@rediffmail.com)  
Web: [stntechnologies.com](http://stntechnologies.com)

## South Delhi Computer Centre

SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT) from KARNATAKA STATE OPEN UNIVERSITY.  
**New Delhi**  
Tel: 011-26183327, Fax: 011-26143642  
Email: [southdelhicomputercentre@gmail.com](mailto:southdelhicomputercentre@gmail.com), [southdelhicomputercentre@hotmail.com](mailto:southdelhicomputercentre@hotmail.com)  
Web: [www.itwhizkid.com](http://www.itwhizkid.com)  
[www.itwhizkid.org](http://www.itwhizkid.org)

## Ssystems Quest

Making Tomorrow's professionals TODAY  
**Bangalore**  
Tel: 080-41301814  
Email: [directorv@ssystemsquest.com](mailto:directorv@ssystemsquest.com)  
Web: [www.ssystemsquest.com](http://www.ssystemsquest.com)

## Trimax FuturePerfect

A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.  
**Mumbai**  
Tel: 022-40681313, Mobile: 09987705638  
Fax: 022-40681001  
Email: [futureperfect@trimax.in](mailto:futureperfect@trimax.in)  
Web: [www.trimax.in](http://www.trimax.in)

## Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award "REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.  
**Mumbai**  
Tel: 022-26285066/6701  
Email: [vibrant@vsnl.net](mailto:vibrant@vsnl.net)  
Web: [www.vibrantcomputers.com](http://www.vibrantcomputers.com)

## Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses adn authorised testing center of Prometric, Vue and Red Hat.  
**Mumbai**  
Tel: 022-67669217  
Email: [unmesh.raote@ultramaxit.com](mailto:unmesh.raote@ultramaxit.com)  
Web: [www.ultramaxit.com](http://www.ultramaxit.com)

## Yash Infotech

Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.  
**Lucknow**  
Tel: 0522-4043386, Fax: 0522-4043386  
Email: [yashinfotech.lko@gmail.com](mailto:yashinfotech.lko@gmail.com)

Want to  
register  
your  
organisation  
in FOSS  
Yellow  
Pages For

**FREE**

Call: Dhiraj (Delhi) 09811206582

Somaiah (Bangalore) 09986075717

or mail: [dhiraj.khare@efyindia.com](mailto:dhiraj.khare@efyindia.com)

or [somaiah.km@efyindia.com](mailto:somaiah.km@efyindia.com)

\*Offer for limited period.

To advertise in this section, please contact  
Somaiah (Bangalore) 09986075717  
Dhiraj (Delhi) 09811206582

# Which one is your career **goal?**



medical



engineering



computer software



computer hardware



hotel management



clinical research

## We help you **decide**

### Hundreds of Educational Institutions Offering Professional Courses & Training Under A Single Roof

#### **EduTech's Highlights :**

- \* An event totally focused on professional education
- \* A one-stop facility for students seeking admission into professional courses
- \* Counselling by industry top-shots (HR and recruitments)
- \* Major plan for advertising and promotion of EduTech Expo 2009
- \* Powered by EFY Group's vision to boost professional education
- \* Courses catering school pass-outs & fresh graduates



EFY Enterprises Pvt Ltd,  
D-87/1, Okhla Indl. Area, Phase-I, New Delhi-110020;  
Ph: 011-26810601-03; Fax: 011-26817563

# **EDUTECH EXPO 2009**

13<sup>th</sup> - 14<sup>th</sup> JUNE 2009  
PRAGATI MAIDAN, NEW DELHI

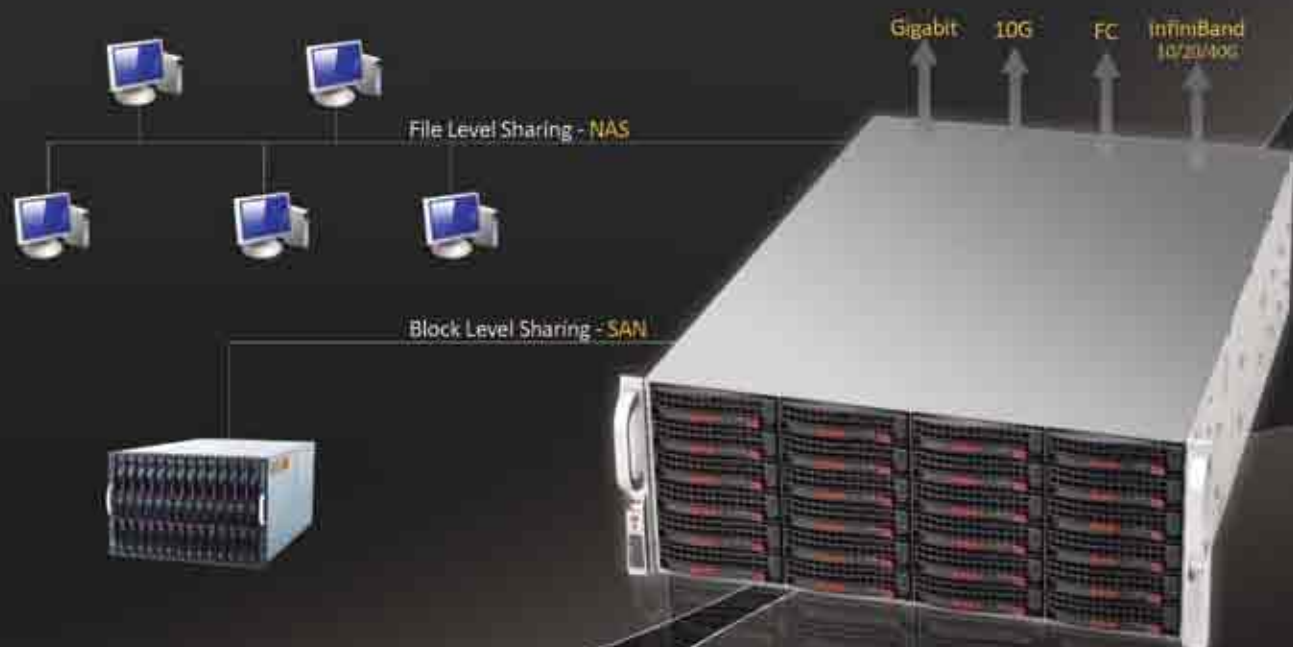






**Tyrone**  
www.tyrone-systems.com

A solution **FLEXIBLE** enough  
to **ADAPT** to your requirements



**UP TO 2GB/s**  
**UP TO 576 TB**  
**SRP Support**

## Opslag FS2

**Unified NAS & SAN**

**Flexible Host Interface**

Supports Gigabit Ethernet, 10GB Ethernet,  
SDR/DDR/QDR (10Gbps/20Gbps/40Gbps) as individual connectivity or in combination

**Scalable Storage**

Supports up to 576 TB from few TB's using Tyrone's expandable solutions

**Redundancy/ Fault-tolerance**

Power-supply, Host Interface failover, RAID 6 (various level of RAID supported), Multiple Snapshots with scheduling, Volume/  
Share Replication, Remote Replication (Block level replication to a 2nd unit with failover/high availability)

**Supports Virtual Tape Library**

**Block-Storage/SAN Protocols**

ISCSI target, Fibre-Channel target, SRP target

**Supported Clients**

Windows, Linux, FreeBSD\*, Solaris\*, Unix\* (\*no SRP support)

**Netweb**  
TECHNOLOGIES  
SDI INTELIGENT Company

Corporate Office  
B-1/ A25, 2nd Floor  
Mohan Co-Operative Industrial Estate,  
Mathura Road, New Delhi - 110044

New Delhi : +91 11 29942258/29942257  
Bangalore : +91 80 41146565/22861407  
Chennai : +91 44 43531009/26631975  
Kolkata : +91 33 23232055/23232056

Website : [www.netwebindia.com](http://www.netwebindia.com)  
[www.tyrone-systems.com](http://www.tyrone-systems.com)  
Email : [info@netwebindia.com](mailto:info@netwebindia.com)  
[netweb@vsnl.com](mailto:netweb@vsnl.com)